arXiv:1402.7001v1 [cs.LG] 27 Feb 2014

Marginalizing Corrupted Features

Laurens van der Maaten LVDMAATEN@GMAIL.COM
Pattern Recognition and Bioinformatics Group
Delft University of Technology
Mekelweg 4, 2628 CD Delft, The Netherlands

Minmin Chen* M.CHEN@CRITEO.COM
Criteo
Palo Alto, CA 94301, USA

Stephen Tyree SWTYREE@WUSTL.EDU
Kilian Q. Weinberger KILIAN@WUSTL.EDU
Department of Computer Science and Engineering
Washington University in St. Louis
St. Louis, MO 63130, USA

Abstract

The goal of machine learning is to develop predictors that generalize well to test data. Ideally, this is achieved by training on an almost infinitely large training data set that captures all variations in the data distribution. In practical learning settings, however, we do not have infinite data and our predictors may overfit. Overfitting may be combatted, for example, by adding a regularizer to the training objective or by defining a prior over the model parameters and performing Bayesian inference. In this paper, we propose a third, alternative approach to combat overfitting: we extend the training set with infinitely many artificial training examples that are obtained by corrupting the original training data. We show that this approach is practical and efficient for a range of predictors and corruption models. Our approach, called *marginalized corrupted features* (MCF), trains robust predictors by minimizing the expected value of the loss function under the corruption model. We show empirically on a variety of data sets that MCF classifiers can be trained efficiently, may generalize substantially better to test data, and are also more robust to feature deletion at test time.

1. Introduction

Dealing with overfitting is one of the key problems one encounters when training machine-learning models. When a learner overfits, it performs substantially better on its training data than on held-out test data. An overfitted learner relates the label or target value with patterns within the noise instead of within the data distribution. Three approaches are commonly used to combat overfitting: 1. *early-stopping* techniques monitor the performance of the model on a validation set during training, and stop the learning as soon as the validation performance deteriorates; 2. *regularization* techniques encourage the learning to find "simple" models by penalizing "complex" models, *e.g.*, models with large parameter values; and 3. *Bayesian* techniques define a prior distribution over models that favor simple models, and perform predictions by averaging over the model posterior. Regularization and

Bayesian techniques are intimately related because regularization techniques can often be viewed as introducing a prior over models and performing maximum-likelihood estimation.

In this paper, we propose a new alternative to counter overfitting. Instead of requiring the user to define prior distributions or regularizers over *model parameters*, which can be very counter intuitive, we focus on corruptions of the *data*. Our approach is based on the observation that overfitting would completely disappear if we were to train our models on infinite data drawn from the data distribution $\mathcal{P}$. In such a hypothetical scenario, it is impossible to overfit models; and even high-variance models, such as the nearest neighbor classifier (Cover and Hart, 1967), become close to optimal (*viz.* the error of a nearest neighbor classifier is twice the Bayes error). Unfortunately, a learning scenario in which we only obtain a finite training set is more realistic; here, some variations in the data distribution will not be captured and the learned model performs worse at test time than during training.

In many learning scenarios, we may however have some additional knowledge about the data distribution: we might know that certain *corruptions* of data instances do not affect their conditional label distributions. As an example, deleting a few words in a text document rarely changes its topic. With this prior knowledge, we can corrupt existing data to generate new artificial instances that resemble those sampled from the actual data distribution. In fact, we will *corrupt* the existing *finite* training examples with a fixed corrupting distribution to construct an *infinite* training set. We show that for a wide range of learning models and noise distributions, it is practical to train models on such an infinite, augmented training set. We refer to the resulting framework as learning with *marginalized corrupted features* (MCF). So instead of approximating the exact statistics of $\mathcal{P}$ with *finite* data, MCF learns from a slightly modified data distribution $\mathcal{P}'$ with *infinite* data.

Burges and Schölkopf (1997) explicitly augment the training set with additional examples that are corrupted through similar transformations. Although the simplicity of such an approach is appealing, it lacks elegance and the computational cost of processing the additional corrupted training examples is prohibitive for most real-world problems. In contrast, we show that it is efficient to train predictors on an infinite amount of corrupted copies of the training data by marginalizing out the corrupting distribution. In particular, we focus on empirical risk minimization and derive analytical solutions for the expected loss under a large family of corrupting distributions for quadratic and exponential loss functions. This allows us to minimize the expected loss in computational time linear in the number of training examples. For logistic loss functions, which are used in many probabilistic models, we derive practical approximations and upper bounds for the expected value of the loss under the corrupting distribution.

Our augmented data distribution $\mathcal{P}'$ is constructed using a simple stochastic rule: pick one of the finite training examples uniformly at random and transform it with some predefined corrupting distribution. Many corrupting distributions are possible, but in this paper we focus on 1. Gaussian corruption, 2. Poisson corruption, and 3. (unbiased) blankout or dropout corruption (random deletion of features). The Gaussian corruption model is mainly of interest for continuous-valued data sets; special cases of MCF with Gaussian corruption have already been studied in the context of Parzen density estimation (Parzen, 1962) and in the context of vicinal risk minimization (Chapelle et al., 2000). The Poisson corruption model is of interest when the data comprises count vectors, *e.g.*,

in document classification. Poisson corruption is particularly appealing as it introduces *no additional hyper-parameters* and, in our results, improves the test accuracy on almost all data sets that comprise count data. The blankout corruption model is of interest in data sets with heavy-tailed feature distributions, such as filter responses, and in settings where blankout noise is a known source of variance in the original data distribution $\mathcal{P}$—possibly unobserved in the training data. This happens, *e.g.*, in document classification from term-frequency vectors: a big portion of each document (especially after stop-word removal) is sampled from the tail of the power-law distribution, for which the blankout noise model is a surprisingly good approximation: it models the common case that some of the words related to the class of the document are missing, *e.g.* because the author used other synonyms. Blankout corruption is also of interest in the "nightmare at test time" scenario (Globerson and Roweis, 2006) in which some of the features are deleted during testing, *e.g.*, due to sensors failures or because the feature computation exceeds a time budget, and we wish to make predictions that are robust to this feature deletion. Different from previous work (Bhattacharyya et al., 2004; Globerson and Roweis, 2006; Shivaswamy et al., 2006; Trafalis and Gilbert, 2007), our setting focuses on random feature removal and not on the (in practice much rarer) worst-case adversarial setting.

An earlier short paper on this work (van der Maaten et al., 2013) already introduces the MCF learning framework, but the present paper provides a large amount of additional details, analysis, and experiments. The paper makes the following contributions: 1. we introduce learning with marginalized corrupted features, a framework that trains robust classifiers by marginalizing out all possible feature corruptions from a pre-defined distribution; 2. we derive plug-in solutions for the quadratic, exponential, and logistic loss functions for a range of corrupting distributions, which can be incorporated into learning algorithms out-of-the-box; and 3. on several real-world data sets, we show that training with MCF may lead to more robust classifiers and may substantially decrease generalization errors in various learning settings. Given the simplicity and elegance of MCF, it could become a valuable alternative to the common l_1 or l_2-norm regularizers.

The outline of the remainder of this paper is as follows. In Section 2, we discuss prior work that is related to the work presented in this paper. In Section 3, we introduce marginalized corrupted features (MCF) regularization and we derive MCF variants of quadratic and exponential loss. In Section 4, we present approximations and upper bounds that allow MCF to be used with logistic loss functions. Section 5 presents an extensive experimental evaluation in which we compare MCF-regularized predictors with standard predictors in various document and image classification problems, in a classification problem from bioinformatics, and in the "nightmare at test time scenario". Section 6 concludes the paper with a discussion of the results presented in the paper as well as directions for future work.

2. Related work

There are two motivations for training a classifier with MCF: 1. to reduce the effects of overfitting during *training*; and 2. to combat the effect of data corruption during *testing*. Both motivations connect with different lines of prior work, and in this section we discuss both of them.

2.1 Corruption during training

The initial publication by Burges and Schölkopf (1997) has inspired various lines of work that *explicitly* corrupt training data during training. Most prominently, Vincent et al. (2008, 2010) propose to randomly blank out features in the training data that is used as input to autoencoders, whilst leaving the desired output (the original training data) unaltered. The resulting denoising autoencoder model is now commonly used as a building block in deep learning (Maillet et al., 2009; Glorot et al., 2011; Mesnil et al., 2011), and blankout corruption is also increasingly applied on the hidden units of neural networks as a form of regularization (Hinton et al., 2012). Since autoencoders are non-linear, the marginalization over the corrupting distribution cannot be performed analytically in such models. Linear denoising autoencoders (Chen et al., 2012) can be viewed as a special case of MCF that aim to minimize the expected value of the reconstruction error under blankout corruption—and which are stacked in multiple layers. Herbrich and Graepel (2004) propose an elegant generalization of SVMs that learns to be invariant to polynomial input transformations via a semi-definite programming formulation.

In online learning and bandit problems, various studies have shown that it is possible to learn from data that are subject to a (possibly unknown) corrupting distribution (Flaxman et al., 2005; Abernethy et al., 2008; Cesa-Bianchi et al., 2010). For instance, Cesa-Bianchi et al. (2011) show that noise in the data does not affect the convergence rate of online learners. A similar regret bound is proven by Rostamizadeh et al. (2011) for online learning under the presence of randomly missing features (*i.e.* blankout corruption). These studies differ from our work in that they show that (unknown) corruptions in the data do not *impair* learning performance too much. By contrast, our work shows that adding corruptions to the data can actually *improve* learning performance.

2.2 Corruption during testing

Several prior studies consider *implicit* approaches for data sets that are subject to corruptions during *test time* (this is also known as "the nightmare at test-time" scenario). Most of these studies propose to minimize the loss under an adversarial worst-case scenario. In particular, Globerson and Roweis (2006) propose a minimax-formulation in which the loss is minimized assuming maximum "damage" through corruption. By contrast, Dekel and Shamir (2008) propose a linear-programming formulation that minimizes an approximation to the same quantity for margin-based predictors. Other studies (Bhattacharyya et al., 2004; Shivaswamy et al., 2006; Trafalis and Gilbert, 2007; Xu et al., 2009) also use minimax approaches that minimize losses under a worst-case scenario, but they focus on a somewhat different corruption model that adds a uniformly drawn constant, within a fixed interval, to each feature. Livni et al. (2012) consider a scenario in which an adversarial chooses the corrupting distribution from the set of all distributions with a pre-specified variance as to increase the expected loss under the corruption as much as possible. Chechik et al. (2008) propose an algorithm that maximizes the margin in the subspace of the observed features for each training instance to be robust against random feature deletion. Teo et al. (2008) generalize the worst-case scenario from simple feature corruption to obtain invariances to transformations such as image rotations or translations. Their framework incorporates sev-

eral prior formulations on learning with invariants as special cases, most prominently, the formulation of Herbrich and Graepel (2004).

Such previous work differs from our approach in that it focuses on adversarial scenarios and does not consider the corruption analytically in expectation. Further, it mostly focuses on a very specific corruption model. In particular, existing approaches suffer from three main disadvantages: 1. they are complex and computationally expensive; 2. they minimize the loss of a worst-case scenario that is very unlikely to be encountered in practice; and 3. they do not provide a flexible framework that can be used with a variety of models and corrupting distributions. By contrast, we propose a much simpler approach that scales linearly in the number of training samples, that considers an average-case instead of a worst-case scenario, and that can readily be used with a variety of loss functions and a large family of corruption models.

3. Learning with Marginalized Corrupted Features (MCF)

To derive the *marginalized corrupted features* (MCF) framework, we start by defining a corrupting distribution that specifies how training observations $\mathbf{x}$ are transformed into corrupted versions $\tilde{\mathbf{x}}$. Throughout the paper, we assume that the corrupting distribution factorizes over dimensions and that each individual distribution P_E is a member of the natural exponential family. (We will see later that for MCF with quadratic loss functions, these assumptions may be relaxed.) Specifically, we assume a corrupting distribution of the form:

$$p(\tilde{\mathbf{x}}|\mathbf{x}) = \prod_{d=1}^{D} P_E(\tilde{x}_d|x_d;\eta_d), \tag{1}$$

where η_d represents user-defined hyperparameters of the corrupting distribution on dimension d. Corrupting distributions of interest, for P_E, include: 1. independent salt or "blank-out" (or dropout) noise in which the d-th feature is randomly set to zero with probability q_d; 2. bit-swap noise in which the value of the d-th bit is randomly swapped with probability q_d; 3. independent Gaussian noise on the d-th feature with variance σ_d^2; 4. independent Laplace noise on the d-th feature with variance $2\lambda_d^2$; and 5. independent Poisson corruption in which the d-th feature is used as the rate of the Poisson distribution. Note that the definition in (1) allows for different features to use arbitrary different corrupting distributions. In the remainder of this paper, we propose an efficient way to train a classifier on corrupted inputs $\tilde{\mathbf{x}}$, sampled from the distribution $p(\tilde{\mathbf{x}}|\mathbf{x})\mathcal{P}(\mathbf{x})$. There are at least two motivations for doing so:

1. Nightmare at test-time. In the *nightmare at test-time* scenario, introduced by Globerson and Roweis (2006), there is an expectation that corruption will appear during test-time. Although the training data is sampled from $\mathcal{P}(\mathbf{x})$, the test data will be sampled from the distribution $p(\tilde{\mathbf{x}}|\mathbf{x})\mathcal{P}(\mathbf{x})$. Training the classifier on the corrupted data corrects this distribution drift. This scenario appears, for example, in the context of search engines. A feature can be reliably collected for the training set, but during testing it is dropped if its computation time exceeds a pre-specified time limit. The frequency of these feature drop-outs can be measured and incorporated into training with the appropriate dropout

distribution. Similarly, one can imagine scenarios (*e.g.* in robotics or sensor networks) where features correspond to unreliable sensor readings.

2. Regularization. The more common motivation is improved generalization. Here, the corrupting distribution is meant to capture some of the variance inherent in $\mathcal{P}(\mathbf{x})$. In this setting, one must guarantee that the classifier is still applicable for the test case, where data is sampled directly from $\mathcal{P}$. More precisely, it must be *unbiased*. Let $h(\mathbf{x})$ be the classification function, we therefore require of the corrupting distribution that

$$\mathbb{E}[h(\tilde{\mathbf{x}})]_{p(\tilde{\mathbf{x}}|\mathbf{x})} = h(\mathbf{x}). \tag{2}$$

In other words, corrupting an input should not change its expected prediction and a test input obtains the expected prediction that the same input would obtain during training. In the linear case, the classifier is parameterized by a vector $\mathbf{w}$. For regression, the prediction function is defined as $h(\mathbf{x}) = \mathbf{w}^\top \mathbf{x}$ and we require $\mathbb{E}[\tilde{\mathbf{x}}] = \mathbf{x}$. For classification, it is $h(\mathbf{x}) = \text{sign}(\mathbf{w}^\top \mathbf{x})$, which necessitates the slightly weaker condition $\mathbb{E}[\tilde{\mathbf{x}}] = s(\mathbf{x})\mathbf{x}$ for any bounded $s(\mathbf{x}) > 0$. Table 1 showcases several biased and unbiased distributions.

Assume we are provided with a training data set $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$ and a loss function $L(\mathbf{x}, y; \Theta)$, with model parameters Θ. A simple approach to approximately learn from the distribution $p(\tilde{\mathbf{x}}|\mathbf{x})\mathcal{P}(\mathbf{x})$ is to follow the spirit of Burges and Schölkopf (1997) and corrupt each training sample M times, following (1). For each $\mathbf{x}_n \in \mathcal{D}$, this results in corresponding corrupted observations $\tilde{\mathbf{x}}_{nm}$ (with $m = 1, \ldots, M$) and leads to the construction of a new data set $\tilde{\mathcal{D}}$ of size $|\tilde{\mathcal{D}}| = MN$. This new data set can be used for training in an empirical risk minimization framework by minimizing the surrogate loss function:

$$\mathcal{L}(\tilde{\mathcal{D}}; \Theta) = \frac{1}{N} \sum_{n=1}^{N} \frac{1}{M} \sum_{m=1}^{M} L(\tilde{\mathbf{x}}_{nm}, y_n; \Theta), \tag{3}$$

with $\tilde{\mathbf{x}}_{nm} \sim p(\tilde{\mathbf{x}}_{nm}|\mathbf{x}_n)$. Such approaches have recently become popular, in particular, in the deep-learning community as a way to regularize deep networks (Hinton et al., 2012).

Although approaches that explicitly corrupt the training data are effective, they lack elegance and come with high computational costs: the minimization of $\mathcal{L}(\tilde{\mathcal{D}}; \Theta)$ scales linearly in the number of corrupted observations, *i.e.* it scales as $\mathcal{O}(NM)$. It is, however, of interest to consider the limiting case in which $M \to \infty$. In this case, we can apply the *weak law of large numbers* and rewrite $\frac{1}{M} \sum_{m=1}^{M} L(\tilde{\mathbf{x}}_m, y_m; \Theta)$ as its expectation (Duda et al., 2001, §2.10.2) to obtain the following surrogate loss:

$$\mathcal{L}(\mathcal{D}; \Theta) = \frac{1}{N} \sum_{n=1}^{N} \mathbb{E}[L(\tilde{\mathbf{x}}_n, y_n; \Theta)]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)}. \tag{4}$$

Minimizing the expected value of the loss under the corruption model leads to a new approach for training predictors that we refer to as learning with *marginalized corrupted features* (MCF). MCF may lead to algorithms with an $\mathcal{O}(N)$ training complexity in situations in which the expectation in (4) is tractable, which is indeed the case for commonly used loss functions $L(\tilde{\mathbf{x}}_n, y_n; \Theta)$ and corrupting distributions $p(\tilde{\mathbf{x}}|\mathbf{x})$.

In the following, we show that for linear predictors that employ a quadratic or exponential loss function, the required expectations under $p(\tilde{\mathbf{x}}|\mathbf{x})$ in (4) can be computed analytically for all corrupting distributions in the natural exponential family. For linear predictors with logistic loss we derive practical approximations and upper bounds for the expected loss under $p(\tilde{\mathbf{x}}|\mathbf{x})$, which may serve as surrogate loss functions.

3.1 Quadratic loss

Assuming a linear model parametrized by vector $\mathbf{w}$ and a target variable y (for regression, y is continuous; for binary classification, $y \in \{-1, +1\}$), the expected value of the quadratic loss under corrupting distribution $p(\tilde{\mathbf{x}}|\mathbf{x})$ is given by:

$$\begin{aligned}
\mathcal{L}(\mathcal{D}; \mathbf{w}) &= \frac{1}{N} \sum_{n=1}^{N} \mathbb{E} \left[\left(\mathbf{w}^\top \tilde{\mathbf{x}}_n - y_n \right)^2 \right]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} \\
&= \frac{1}{N} \sum_{n=1}^{N} \mathbb{E} \left[\mathbf{w}^\top \tilde{\mathbf{x}}_n \tilde{\mathbf{x}}_n^\top \mathbf{w} - 2 y_n \mathbf{w}^\top \tilde{\mathbf{x}}_n + y_n^2 \right]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} \\
&= \mathbf{w}^\top \left(\frac{1}{N} \sum_{n=1}^{N} \left[\mathbb{E}[\tilde{\mathbf{x}}_n] \mathbb{E}[\tilde{\mathbf{x}}_n]^\top + \operatorname{diag}\left(\mathbb{V}[\tilde{\mathbf{x}}_n] \right) \right] \right) \mathbf{w} - 2 \left(\frac{1}{N} \sum_{n=1}^{N} y_n \mathbb{E}\left[\tilde{\mathbf{x}}_n \right] \right)^\top \mathbf{w} + 1.
\end{aligned} \tag{5}$$

Herein, $\mathbb{V}[x]$ denotes the variance of x, and all expectations are under $p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)$. Note that irrespective of what corruption model is used, (5) is quadratic in $\mathbf{w}$, and hence convex. The optimal solution $\mathbf{w}^*$ can be computed with a variation of the ordinary least squares closed-form solution:

$$\mathbf{w}^* = \left(\sum_{n=1}^{N} \left[\mathbb{E}[\tilde{\mathbf{x}}_n] \mathbb{E}[\tilde{\mathbf{x}}_n]^\top + \operatorname{diag}\left(\mathbb{V}[\tilde{\mathbf{x}}_n] \right) \right] \right)^{-1} \left(\sum_{n=1}^{N} y_n \mathbb{E}\left[\tilde{\mathbf{x}}_n \right] \right). \tag{6}$$

Hence, to minimize the expected quadratic loss under the corruption model, we only need to compute the mean and variance of the corrupting distribution, which is practical for all exponential-family distributions (as well as for certain distributions outside of the exponential family, such as Student-t distributions with more than two degrees of freedom). Table 1 gives an overview of the required mean and variance for six corrupting distributions of interest.

Special case. An interesting setting of MCF with quadratic loss occurs when the corrupting distribution $p(\tilde{\mathbf{x}}|\mathbf{x})$ is an isotropic Gaussian distribution with mean $\mathbf{x}$ and variance $\sigma^2\mathbf{I}$. For such a Gaussian corruption model, we obtain the standard l_2-regularized quadratic loss with regularization parameter σ^2 that is used in ridge regression as special case (Chapelle et al., 2000):

$$\mathcal{L}(\mathcal{D}; \mathbf{w}) = \mathbf{w}^\top \left(\frac{1}{N} \sum_{n=1}^{N} \mathbf{x}_n \mathbf{x}_n^\top \right) \mathbf{w} - 2 \left(\frac{1}{N} \sum_{n=1}^{N} y_n \mathbf{x}_n \right)^\top \mathbf{w} + \sigma^2 \mathbf{w}^\top \mathbf{w} + 1. \tag{7}$$

Perhaps surprisingly, using MCF with Laplace corruption also leads to ridge regression with the regularization parameter taking value $2\lambda^2$. Indeed, ridge regression arises for every

	Name	PDF	$\mathbb{E}[\tilde{\boldsymbol{x}}_{\boldsymbol{nd}}]_{p(\tilde{\boldsymbol{x}}_{\boldsymbol{nd}}\mid\boldsymbol{x}_{\boldsymbol{nd}})}$	$\boldsymbol{V}[\tilde{\boldsymbol{x}}_{\boldsymbol{nd}}]_{p(\tilde{\boldsymbol{x}}_{\boldsymbol{nd}}\mid\boldsymbol{x}_{\boldsymbol{nd}})}$
biased	Dropout	$p(\tilde{x}_{nd}=0)=q_d$ $p(\tilde{x}_{nd}=x_{nd})=1-q_d$	$(1-q_d)x_{nd}$	$q_d(1-q_d)x_{nd}^2$
biased	Bit-swap	$p(\tilde{x}_{nd}=1-x_{nd})=q_d$ $p(\tilde{x}_{nd}=x_{nd})=1-q_d$	$q^{(1-x_{nd})}(1-q)^{x_{nd}}$	$\left(1-q^{(1-x_{nd})}(1-q)^{x_{nd}}\right)$ $q^{(1-x_{nd})}(1-q)^{x_{nd}}$
unbiased	Blankout	$p(\tilde{x}_{nd}=0)=q_d$ $p(\tilde{x}_{nd}=\frac{1}{1-q_d}x_{nd})=1-q_d$	x_{nd}	$\frac{q_d}{1-q_d}x_{nd}^2$
unbiased	Gaussian	$p(\tilde{x}_{nd}\mid x_{nd})=\mathcal{N}(\tilde{x}_{nd}\mid x_{nd},\sigma^2)$	x_{nd}	σ^2
unbiased	Laplace	$p(\tilde{x}_{nd}\mid x_{nd})=Lap(\tilde{x}_{nd}\mid x_{nd},\lambda)$	x_{nd}	$2\lambda^2$
unbiased	Poisson	$p(\tilde{x}_{nd}\mid x_{nd})=Poisson(\tilde{x}_{nd}\mid x_{nd})$	x_{nd}	x_{nd}

Table 1: The probability density function (PDF), mean, and variance of corrupting distributions of interest. These quantities can be plugged into Eq. (5) to obtain the expected value under the corrupting distribution of the quadratic loss. For disambiguation, we refer to the unbiased version of feature dropout as blankout corruption.

unbiased corrupting distribution whose variance is not data-dependent. For corrupting distributions whose variance is data-dependent, other regularizers are obtained (see below).

Blankout corruption. We study the use of blankout corruption in MCF quadratic loss in more detail here to show what kind of regularizers it produces. By plugging the relevant quantities from Table 1 into (6), we obtain the following solution for $\mathbf{w}^*$ with MCF blankout (assuming $\forall d : q_d = q$):

$$\mathbf{w}^* = \left(\sum_{n=1}^{N}\mathbf{x}_n\mathbf{x}_n^\top + q(1-q)\mathbf{S}\right)^{-1}\left(\sum_{n=1}^{N}\mathbf{x}_n y_n\right), \text{ where: } \mathbf{S} = \left(\sum_{n=1}^{N}\mathbf{x}_n\mathbf{x}_n^\top\right)\circ\mathbf{I} \tag{8}$$

where $\circ$ denotes an element-wise Hadamard product. In words, the matrix $\mathbf{S}$ consists of only the diagonal of the Gram matrix. Consequently, MCF blankout corruption has the effect of multiplying the diagonal of the Gram matrix by a factor $(1+q-q^2)$. This makes the MCF regularizer *data-dependent*: the regularization is stronger on dimensions that on average have a larger norm. This may be desirable in situations in which the data dimensions live on different scales as happens, for example, in word-count features. This is in contrast to the *data-independent* l_2-regularizer, that simply adds a fixed value λ to the diagonal of the scatter matrix—treating all features alike. Similar data-dependent regularizers have previously been studied by Grave et al. (2011).

Poisson corruption. Using Poisson corruption in MCF quadratic loss leads to:

$$\mathbf{w}^* = \left(\sum_{n=1}^{N}\mathbf{x}_n\mathbf{x}_n^\top + \operatorname{diag}\left(\sum_{n=1}^{N}\mathbf{x}_n\right)\right)^{-1}\left(\sum_{n=1}^{N}\mathbf{x}_n y_n\right). \tag{9}$$

Poisson corruption thus leads to a *parameterless, additive, data-dependent* regularizer that adds larger values to the diagonal of the Gram matrix for features that occur more on

average. (Note that Poisson corruption can only be used on non-negative data, $\forall n, d : x_{nd} \geq 0$, so the contribution of the regularizer is always non-negative.) The main difference between blankout and Poisson corruption is in how strongly features with large values are regularized; the multiplicative blankout corruption regularizes common features more than the additive Poisson corruption.

Leave-one-out errors. A nice property of linear models employing quadratic loss is that they allow us to compute the leave-one-out error of the model from only the training predictions $\hat{y}_n$ (Allen, 1974). We show that this property still holds for MCF quadratic losses, irrespective of the corrupting distribution that is used. Defining the target vector $\mathbf{y} = [y_1, \ldots, y_N]^\top$ and the data matrix $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_N]$, we can rewrite the prediction vector $\hat{\mathbf{y}} = [\hat{y}_1, \ldots, \hat{y}_N]$ as:

$$\hat{\mathbf{y}} = \mathbf{X}^\top \mathbf{w}^* = \mathbf{X}^\top \left(\sum_{n=1}^{N} \left[\mathbb{E}[\tilde{\mathbf{x}}_n] \mathbb{E}[\tilde{\mathbf{x}}_n]^\top + \operatorname{diag}\left(\mathbb{V}[\tilde{\mathbf{x}}_n] \right) \right] \right)^{-1} \left(\sum_{n=1}^{N} \mathbb{E}\left[\tilde{\mathbf{x}}_n \right] y_n \right) = \mathbf{H}\mathbf{y}, \quad (10)$$

where $\mathbf{H}$ is often referred to as the hat matrix. If a single target value y_n is replaced by $\hat{y}_n^{-n}$, the prediction made by the model trained on all data except example n, the hat matrix $\mathbf{H}$ does not change because it does not depend on $\mathbf{y}$. If we construct a new target vector $\mathbf{z}$ with this replacement in it, $\mathbf{z} = [y_1, \ldots, y_{n-1}, \hat{y}_n^{-n}, y_{n+1}, \ldots, y_N]$, and use the simple fact that $\hat{y}_n^{-n} = \sum_{m=1}^{N} H_{nm} z_m$ (see Allen (1974) for details) we obtain:

$$\hat{y}_n - \hat{y}_n^{-n} = \sum_{m=1}^{N} H_{nm} y_m - \sum_{m=1}^{N} H_{nm} z_m = H_{nn} y_n - H_{nn} \hat{y}_n^{-n}. \quad (11)$$

Using this equality, we can compute $\hat{y}_n^{-n}$ without performing the minimization of the MCF quadratic loss – with example n left out – over $\mathbf{w}$:

$$\hat{y}_n^{-n} = \frac{H_{nn} y_n - \hat{y}_n}{H_{nn} - 1}. \quad (12)$$

This allows us to express the leave-one-out error LOO of the MCF quadratic loss model as:

$$LOO = \frac{1}{N} \sum_{n=1}^{N} \left(y_n - \hat{y}_n^{-n} \right)^2 = \frac{1}{N} \sum_{n=1}^{N} \left(\frac{y_n - \hat{y}_n}{1 - H_{nn}} \right)^2. \quad (13)$$

3.2 Exponential loss

Whilst quadratic losses are of interest to regression problems, they seem less appropriate for use in classification. Exponential loss is a loss function that is commonly used for classification, *e.g.*, in AdaBoost (Freund and Schapire, 1995). Assuming a label variable $y \in \{0, 1\}$, the expected value of the exponential loss under corruption model $p(\tilde{\mathbf{x}}|\mathbf{x})$ is given by:

$$\begin{aligned} \mathcal{L}(\mathcal{D}; \mathbf{w}) &= \frac{1}{N} \sum_{n=1}^{N} \mathbb{E}\left[\exp\left(-y_n \mathbf{w}^\top \tilde{\mathbf{x}}_n \right) \right]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} \\ &= \frac{1}{N} \sum_{n=1}^{N} \prod_{d=1}^{D} \mathbb{E}\left[\exp\left(-y_n w_d \tilde{x}_{nd} \right) \right]_{p(\tilde{x}_{nd}|x_{nd})}, \end{aligned} \quad (14)$$

Distribution	$\mathbb{E}[\exp(-y_n w_d \tilde{x}_{nd})]_{p(\tilde{x}_{nd}\|x_{nd})}$	$\partial \mathbb{E}[\exp(-y_n w_d \tilde{x}_{nd})]_{p(\tilde{x}_{nd}\|x_{nd})} / \partial w_d$
Blankout noise	$q_d + (1-q_d)\exp(-y_n w_d \frac{1}{1-q_d} x_{nd})$	$-\exp(-y_n w_d \frac{1}{1-q_d} x_{nd}) y_n x_{nd}$
Gaussian noise	$\exp(-y_n w_d x_{nd} + \frac{1}{2}\sigma^2 y_n^2 w_d^2)$	$\exp(-y_n w_d x_{nd} + \frac{1}{2}\sigma^2 y_n^2 w_d^2)(-y_n x_{nd} + \sigma^2 y_n^2 w_d)$
Laplace noise	$(1-\lambda^2 y_n^2 w_d^2)^{-1}\exp(-y_n w_d x_{nd})$	$-\exp(-y_n w_d x_{nd})[(1-\lambda^2 y_n^2 w_d^2)^{-1} y_n x_{nd} + 2(1-\lambda^2 y_n^2 w_d^2)^{-2}\lambda^2 y_n^2 w_d]$
Poisson noise	$\exp(x_{nd}(\exp(-y_n w_d)-1))$	$-\exp(x_{nd}(\exp(-y_n w_d)-1) - y_n w_d) y_n x_{nd}$

Table 2: Moment-generating functions (MGFs) of various corrupting distributions, and the corresponding gradients of these functions with respect to the weights w_d. These quantities can be plugged into equations (14) and (25) to obtain the expected value of the loss (or surrogate) under the corrupting distribution of the exponential and logistic loss functions, respectively.

which we refer to as MCF exponential loss. Note that in the derivation, we used the assumption that the corruption is independent across features. The above equation can be recognized as a product of moment-generating functions $\mathbb{E}[\exp(t_{nd}\tilde{x}_{nd})]$ with $t_{nd} = -y_n w_d$. By definition, the moment-generating function (MGF) can be computed for all corrupting distributions that are member of the natural exponential family. An overview of the moment-generating functions for some corrupting distributions of interest is given in Table 2. Noting that $f(x)$ is convex iff $\forall x : \frac{\partial^2 f(x)}{\partial x^2} \geq 0$ and noting that sums of convex functions are themselves convex, it is straightforward to verify that MCF exponential loss is convex for all corrupting distributions. We also note that for corrupting distributions that do not permit a moment-generating function but that have bounded support, it may be practical to upper bound (14) using Hoeffding's lemma; we leave such investigations to future work.

The derivation of (14) can readily be extended to a multi-class exponential loss with K classes by replacing the weight vector $\mathbf{w}$ by a $D \times K$ weight matrix $\mathbf{W}$, and by replacing the labels y by label vectors $\mathbf{y} = \{1, -\frac{1}{K-1}\}^K$ with $\sum_{k=1}^K y_k = 0$ (Zhu et al., 2006).

Unlike the minimization of MCF quadratic loss, the minimization of MCF exponential losses cannot be done in closed form and needs to be performed using gradient-descent techniques. Motivated by Sha and Pereira (2003), we used an L-BFGS optimizer to minimize MCF exponential loss in this study.

Blankout corruption. As an illustrative example, we work out MCF exponential loss with blankout corruption:

$$\mathcal{L}(\mathcal{D}; \mathbf{w}) = \frac{1}{N}\sum_{n=1}^{N}\prod_{d=1}^{D}\left(q_d + (1-q_d)\exp(-y_n w_d x_{nd})\right). \tag{15}$$

The gradient of this loss function is given by:

$$\frac{\partial \mathcal{L}}{\partial w_d} = -\frac{1}{N}\sum_{n=1}^{N}(1-q_d)\exp(-y_n w_d x_{nd}) y_n x_{nd} \prod_{d' \neq d}\left(q_{d'} + (1-q_{d'})\exp(-y_n w_{d'} x_{nd'})\right). \tag{16}$$

Although the complexity of training with MCF remains $\mathcal{O}(N)$, evaluating this gradient requires more computation than evaluating the gradient of a standard exponential loss: in

practice, the computation of the MCF exponential loss gradient is about twice as expensive when computation from the loss (15) is reused efficiently.

4. MCF for Logistic Loss

In the case of the logistic loss, the solution to the expected loss (4) cannot be computed in closed form. Instead, we derive one approximation and two practical upper bounds for the expected logistic loss, all of which can be used as surrogate loss functions in practice. In this section we follow the framework of Wager et al. (2013), which is clearer than our previously used derivation (van der Maaten et al. (2013)) and leads to a nice interpretation of MCF as an adaptive regularizer.

We will focus on binary classification in our derivations below, but our results can readily be extended to multi-class logistic loss. (In general, this is achieved by redefining the labels y to be label vectors of the form $\mathbf{y} \in \{0,1\}^K$ with $\sum_{k=1}^K y_k = 1$; and by defining the loss as the logarithm of the softmax probability of the correct prediction.)

4.1 MCF Regularizer

Assuming labels $y \in \{-1,+1\}$, we may define the logistic regression model as a generalized linear model, $p(y|x;\mathbf{w}) = \frac{\exp(y\mathbf{w}^\top\mathbf{x})}{\exp(-\mathbf{w}^\top\mathbf{x})+\exp(\mathbf{w}^\top\mathbf{x})}$. The negative log-likelihood, or the *logistic loss*, of an instance pair $(\mathbf{x}, y)$ under this model can be written as:

$$L(\mathbf{x}, y, \mathbf{w}) = -y\mathbf{w}^\top\mathbf{x} + A(\mathbf{w}^\top\mathbf{x}) \text{ with: } A(\mathbf{w}^\top\mathbf{x}) = \log\left[\exp\left(-\mathbf{w}^\top\mathbf{x}\right) + \exp\left(\mathbf{w}^\top\mathbf{x}\right)\right]. \quad (17)$$

Here, $A(\mathbf{w}^\top\mathbf{x})$ is the log-partition function, which is independent of the label y. Noting that we assumed the corruption to be unbiased (*i.e.* $\mathbb{E}[\tilde{\mathbf{x}}] = \mathbf{x}$), we can plug (17) into the expected loss (4) and obtain:

$$\begin{aligned}\mathcal{L}(\mathcal{D};\mathbf{w}) &= -\frac{1}{N}\sum_{n=1}^{N}\mathbb{E}\left[y_n\mathbf{w}^\top\tilde{\mathbf{x}}_n - A(\mathbf{w}^\top\tilde{\mathbf{x}}_n)\right]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} \\ &= \frac{1}{N}\sum_{n=1}^{N}\left(\underbrace{L(\mathbf{x}_n, y_n, \mathbf{w})}_{\text{log-loss}} + \underbrace{\mathbb{E}\left[A(\mathbf{w}^\top\tilde{\mathbf{x}}_n)\right]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} - A(\mathbf{w}^\top\mathbf{x}_n)}_{\text{regularizer}}\right). \quad (18)\end{aligned}$$

Wager et al. (2013) observe that the first part of (18) is the regular logistic loss without MCF, and the second part can be interpreted as an adaptive regularizer that satisfies several appealing properties: 1. it only depends on $\mathbf{w}$ and $\mathbf{x}_n$ but not on the label y; 2. it is always non-negative—which follows directly from the fact that $A(\mathbf{w}^\top\mathbf{x})$ is convex and an application of Jensen's inequality, *i.e.* $\mathbb{E}[A(\mathbf{w}^\top\tilde{\mathbf{x}})] \geq A(\mathbb{E}[\mathbf{w}^\top\tilde{\mathbf{x}}]) = A(\mathbf{w}^\top\mathbf{x})$; and 3. it has an intuitive interpretation, as it minimizes the difference between the expected prediction before and after corruption—in other words, the score that the classifier assigns to an example $\mathbf{x}_n$ should be robust with respect to the corruption distribution.

Unfortunately, $\mathbb{E}[A(\mathbf{w}^\top\tilde{\mathbf{x}})]$ in (18) cannot be expressed in closed form. In the following three sections we therefore derive practical closed-form approximations and upper bounds that can be used to form surrogate loss functions for (18).

4.2 Quadratic Approximation

A quadratic approximation to the expected logistic loss can be obtained by deriving a second-order Taylor expansion of $A(\mathbf{w}^\top \tilde{\mathbf{x}})$ around $\mathbf{w}^\top \mathbf{x}$ and working out the expectation (Wager et al., 2013). The resulting quadratic approximation to the expected logistic loss then becomes:

$$\mathcal{L}(\mathcal{D}; \mathbf{w}) \approx \frac{1}{N} \sum_{n=1}^{N} \left[-y_n \mathbf{w}^\top \mathbf{x}_n + A(\mathbf{w}^\top \mathbf{x}_n) + \frac{1}{2} \frac{\partial^2 A(\mathbf{w}^\top \mathbf{x}_n)}{\partial (\mathbf{w}^\top \mathbf{x}_n)^2} \mathbb{V}[\mathbf{w}^\top \tilde{\mathbf{x}}_n]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} \right] \tag{19}$$

$$= \frac{1}{N} \sum_{n=1}^{N} \left[L(\mathbf{x}_n, y_n, \mathbf{w}) + \frac{1}{2} \frac{\partial^2 A(\mathbf{w}^\top \mathbf{x}_n)}{\partial (\mathbf{w}^\top \mathbf{x}_n)^2} \mathbb{V}[\mathbf{w}^\top \tilde{\mathbf{x}}_n]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} \right]. \tag{20}$$

In this equation, the first-order term has vanished because we assumed the corrupting distribution to be unbiased, $\mathbb{E}[\tilde{\mathbf{x}}] = \mathbf{x}$. The expression $\mathbb{V}[\mathbf{w}^\top \tilde{\mathbf{x}}_n]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)}$ denotes the variance of $\mathbf{w}^\top \tilde{\mathbf{x}}_n$ under the corrupting distribution, which takes the form:

$$\mathbb{V}[\mathbf{w}^\top \tilde{\mathbf{x}}_n]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} = \mathbf{w}^\top \left(\mathbb{E}[\tilde{\mathbf{x}}_n \tilde{\mathbf{x}}_n^\top]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} - \mathbf{x}_n \mathbf{x}_n^\top \right) \mathbf{w}. \tag{21}$$

The second derivative of the log-partition function takes the form:

$$\frac{\partial^2 A(\mathbf{w}^\top \mathbf{x}_n)}{\partial (\mathbf{w}\mathbf{x}_n)^2} = \frac{4 \exp(-2\mathbf{w}^\top \mathbf{x}_n)}{(1 + \exp(-2\mathbf{w}^\top \mathbf{x}_n))^2} = 4\sigma(2\mathbf{w}^\top \mathbf{x}_n)(1 - \sigma(2\mathbf{w}^\top \mathbf{x}_n)) = \mathbb{V}[\hat{y}_n], \tag{22}$$

where $\sigma(z)$ denotes the sigmoid function, $\sigma(z) = \frac{1}{1+\exp(-z)}$, and $\mathbb{V}[\hat{y}_n]$ denotes the variance of the prediction $\hat{y}_n = \text{sign}(\mathbf{w}^\top \mathbf{x}_n)$ under the logistic regression model[1]. Hence, the quadratic approximation (20) can be interpreted as the sum of the standard logistic loss and a data-dependent, additive regularizer that is a product of: 1. the variance of the classifier score under the data corruption and 2. the variance of the prediction under the logistic regression model. Because the regularizer is the product of two variances, it is guaranteed to be non-negative. The additive regularizer vanishes when no corruption is applied on the data, because (21) is zero when the data is not corrupted.

4.3 Jensen's upper bound

As an alternative to the quadratic approximation, we can obtain an effective upper bound on the expected logistic loss under the corrupting distribution using Jensen's inequality, which states that, because $\log(\cdot)$ is a concave function, $\mathbb{E}[\log(z)] \leq \log(\mathbb{E}[z])$. In particular, we obtain the following upper bound on the log-partition function:

$$\mathbb{E}\left[\log\left(\exp\left(-\mathbf{w}^\top \tilde{\mathbf{x}} \right) + \exp\left(\mathbf{w}^\top \tilde{\mathbf{x}} \right) \right) \right] \leq \log\left(\mathbb{E}\left[\exp\left(-\mathbf{w}^\top \tilde{\mathbf{x}} \right) \right] + \mathbb{E}\left[\exp\left(\mathbf{w}^\top \tilde{\mathbf{x}} \right) \right] \right), \tag{23}$$

where all expectations are under the corrupting distribution $p(\tilde{\mathbf{x}}|\mathbf{x})$.

An alternative, somewhat simpler way to express the Jensen's upper bound on $\mathcal{L}(\mathcal{D}; \mathbf{w})$ can be achieved by redefining the labels $y \in \{0, 1\}$, and rephrasing (18) as:

$$\mathcal{L}(\mathcal{D}; \mathbf{w}) = \frac{1}{N} \sum_{n=1}^{N} \mathbb{E}\left[\log\left(1 + \exp\left(-y_n \mathbf{w}^\top \tilde{\mathbf{x}}_n \right) \right) \right]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)}, \tag{24}$$

1. Please note that the mean $\mathbb{E}[\hat{y}] = \frac{\partial A(\mathbf{w}^\top \mathbf{x})}{\partial (\mathbf{w}\mathbf{x})} = 2\sigma(2\mathbf{w}^\top \mathbf{x}) - 1$.

Applying Jensen's inequality for $\log(\cdot)$ to (24), and subsequently following the same steps as in (14), leads to:

$$\mathcal{L}(\mathcal{D};\mathbf{w}) \leq \frac{1}{N}\sum_{n=1}^{N}\log\left(1+\prod_{d=1}^{D}\mathbb{E}\left[\exp\left(-y_n w_d \tilde{x}_{nd}\right)\right]_{p(\tilde{x}_{nd}|x_{nd})}\right). \tag{25}$$

We again recognize a product of MGFs, which can be computed in closed-form for corrupting distributions in the natural exponential family (see Table 2). The upper bound on the expected logistic loss (25) can be rewritten as a log-sum-exp or log-integral-exp function (depending on whether the corruption is discrete or continuous) of an affine function of the parameters $\mathbf{w}$, as a result of which it is guaranteed to be convex (Boyd and Vandenberghe, 2004, §3.1.5). We note here that the evaluation of (25) in a numerically stable way is non-trivial when blankout corruption is used; for details on numerical stability, we refer to Appendix B.

Poisson corruption. As an illustrative example, we show the upper bound of the logistic loss (25), where inputs are corrupted with the Poisson distribution:

$$\mathcal{L}(\mathcal{D};\mathbf{w}) \leq \frac{1}{N}\sum_{n=1}^{N}\log\left(1+\exp\left(\sum_{d=1}^{D}x_{nd}(\exp(-y_n w_d)-1)\right)\right). \tag{26}$$

As with all MCF losses that use Poisson corruption, this regularized loss does not have any additional hyper-parameters. Further, the sum over all features can still be computed efficiently for sparse data by summing only over non-zero entries in $\mathbf{x}_n$. As for MCF exponential loss, we perform the minimization of this surrogate loss with an L-BFGS optimizer. The gradient of (26) is given by:

$$\frac{\partial\mathcal{L}}{\partial w_d} = -\frac{1}{N}\sum_{n=1}^{N}\frac{\exp\left(\sum_{d'\neq d}x_{nd'}(\exp(-y_n w_{d'})-1)\right)}{1+\exp\left(\sum_{d=1}^{D}x_{nd}(\exp(-y_n w_d)-1)\right)}\exp(y_n w_d)y_n x_{nd}. \tag{27}$$

It is fair to say that in practice, the computation of the gradient of this MCF logistic loss takes approximately $10\times$ longer than the computation of a regular logistic loss gradient (due to extra exponentiations and multiplications in the gradient computation).

4.4 Variational bound

Indeed, Jensen's inequality is not the only bound that may be used to derive an upper bound on MCF logistic loss. In particular, we also consider variational bounds of the form:

$$\begin{aligned}
\mathcal{L}(\mathcal{D};\mathbf{w}) &= \frac{1}{N}\sum_{n=1}^{N}\mathbb{E}\left[\log(1+\exp(-y_n\mathbf{w}^\top\tilde{\mathbf{x}}_n))\right]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} \\
&= \min_{\lambda\in[0,1]}\frac{1}{N}\sum_{n=1}^{N}\left(\mathbb{E}\left[-\lambda y_n\mathbf{w}^\top\tilde{\mathbf{x}}_n\right]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)}+\mathbb{E}\left[\log(\exp(\lambda y_n\mathbf{w}^\top\tilde{\mathbf{x}}_n)(1+\exp(-y_n\mathbf{w}^\top\tilde{\mathbf{x}}_n)))\right]\right) \\
&\leq \min_{\lambda\in[0,1]}\frac{1}{N}\sum_{n=1}^{N}\left(\log\left(\mathbb{E}\left[\exp(\lambda y_n\mathbf{w}^\top\tilde{\mathbf{x}}_n)\right]+\mathbb{E}\left[\exp((\lambda-1)y_n\mathbf{w}^\top\tilde{\mathbf{x}}_n)\right]\right)-\lambda y_n\mathbf{w}^\top\mathbf{x}_n\right),
\end{aligned} \tag{28}$$

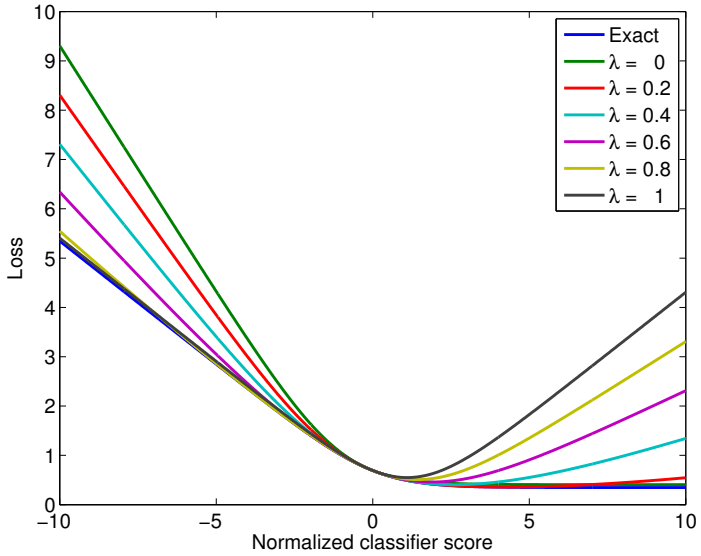

Figure 1: Variational upper bounds on the MCF-blankout logistic loss with $q = 0.5$ for various values of λ. Please note that the Jensen's upper bound of Section 4.3 corresponds to the case $\lambda = 0$.

where we use the assumption that the corrupting distribution is unbiased, $\mathbb{E}[\tilde{\mathbf{x}}] = \mathbf{x}$. The expression above contains two moment-generating functions (with $t = \lambda y_n \mathbf{w}$ and $t = (\lambda - 1) y_n \mathbf{w}$, respectively), which can be evaluated as before. Minimizing (28) over the hyperparameter $\lambda \in [0, 1]$, using a grid or gradient search, leads to a potentially tighter bound than that in (25). Note that the special case $\lambda = 0$ equals (25), as a result of which the alternative bound is always at least as tight. The bounds are illustrated for various values of λ in Figure 1.

4.5 Graphical model interpretation

Finally, minimizing MCF logistic loss can also be interpreted as training the parameters of the simple Bayesian network in Figure 2 via maximum likelihood. In this network, $p(\tilde{x}_d|x_d)$ denotes a corrupting distribution, and $p(y|\tilde{\mathbf{x}})$ comprises a simple linear logistic regressor: $p(y|\tilde{\mathbf{x}}) = 1/(1 + \exp(-y\mathbf{w}^\top \tilde{\mathbf{x}}))$. Marginalizing out the corrupted data $\tilde{\mathbf{x}}$, we obtain:

$$\begin{aligned} p(y_n|\mathbf{x}_n) &= \mathbb{E}\left[\frac{1}{1 + \exp(-y_n \mathbf{w}^\top \tilde{\mathbf{x}}_n)}\right]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)} \\ &\geq \frac{1}{1 + \mathbb{E}\left[\exp(-y_n \mathbf{w}^\top \tilde{\mathbf{x}}_n)\right]_{p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)}}, \end{aligned} \tag{29}$$

where we have used the fact that $\frac{1}{1+z}$ is convex on $[0, \infty)$ and that $\forall z : 1 + \exp(z) > 1$ to lower bound $p(y|\mathbf{x})$ using Jensen's inequality. (Indeed, we could also have applied the quadratic approximation of 4.2 or the variational approximation of 4.4 here.)

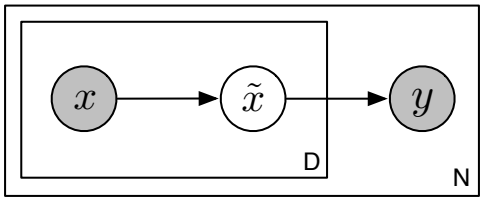

Figure 2: Graphical model interpretation of marginalized corrupted features (MCF). Shaded circles represent observed variables; non-shaded circles represent latent variables; and plates denote independent repetitions. The arrow from x_{nd} to $\tilde{x}_{nd}$ represents the corrupting distribution $p(\tilde{x}_{nd}|x_{nd})$ and the arrow from $\tilde{\mathbf{x}}_n$ to y_n represents the logistic regression model $p(y_n|\tilde{\mathbf{x}}_n)$.

Because the corrupting distribution $p(\tilde{\mathbf{x}}|\mathbf{x})$ factorizes over dimensions, the log-likelihood ℓ of data $\mathcal{D}$ can thus be lower bounded by:

$$\ell(\mathcal{D};\mathbf{w}) \geq -\sum_{n=1}^{N} \log\left(1 + \prod_{d=1}^{D} \mathbb{E}\left[\exp(-y_n w_d \tilde{x}_{nd})\right]_{p(\tilde{x}_{nd}|x_{nd})}\right). \tag{30}$$

It is straightforward to verify that maximizing this lower bound on the log-likelihood is identical to minimizing (25). Further, it it possible to develop a similar graphical model for MCF-quadratic loss by setting $p(y_n|\tilde{\mathbf{x}}_n) = \mathcal{N}(y_n|\mathbf{w}^\top\tilde{\mathbf{x}}, \sigma^2)$.

The graphical model interpretation of MCF also suggests an alternative way to make predictions with models trained using MCF, *viz.* by evaluating (29). In practice, this amounts to generating infinitely many corrupted copies of the test data using $p(\tilde{\mathbf{x}}|\mathbf{x})$ and *averaging over the predictions for all corrupted test points.* While corrupting test data during prediction may sound counterintuitive at first, it does provide an elegant way to remove the restriction for MCF corruption distributions to be unbiased (especially in the case of regression). In classification, corrupting test data provides a natural way to measure the uncertainty of a prediction, even when the original predictor was non-probabilistic (much like conformal prediction; Shafer and Vovk (2008)).

5. Experiments

We evaluate MCF predictors on four tasks: 1. document classification based on word-count features using MCF with blankout and Poisson corruption; 2. image classification based on bag-of-visual-word features using MCF with blankout and Poisson corruption; 3. splice-junction recognition in DNA sequences using blankout and multinomial corruption; and 4. classification of objects in the "nightmare at test time" scenario using MCF with blankout corruption. The four sets of experiments are described separately below. Code to reproduce the results of our experiments is available online[2].

2. A Matlab implementation is available at `http://bit.ly/11bn2GG`.

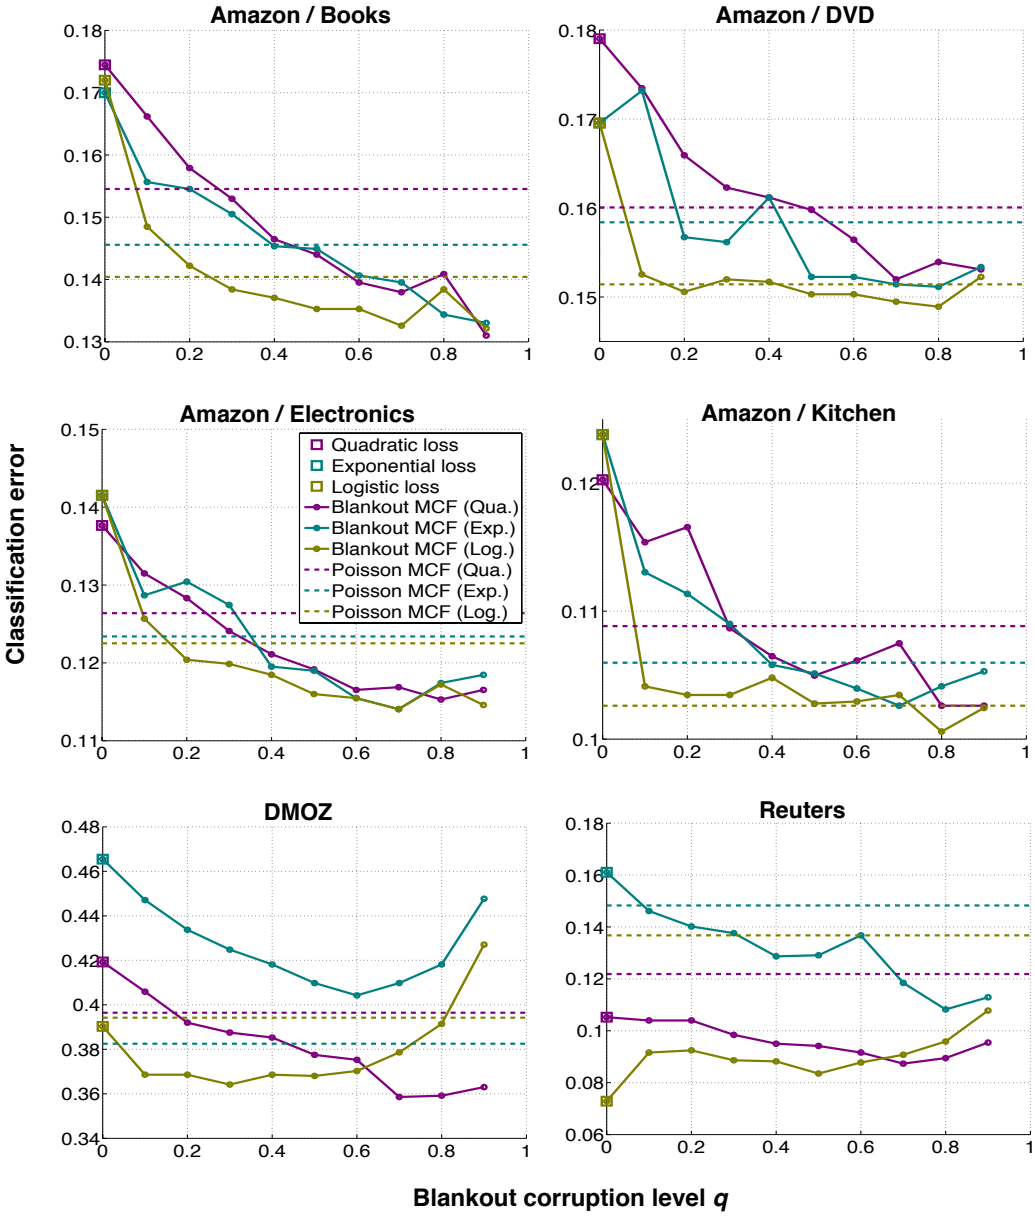

Figure 3: Classification errors of MCF predictors using blankout and Poisson corruption – as a function of the blankout corruption level q – on the Amazon data set for l_2-regularized quadratic, exponential, and quadratic loss functions. Classification errors are represented on the y-axis, whereas the blankout corruption level q is represented on the x-axis. The case of MCF with blankout corruption and $q = 0$ corresponds to a standard l_2-regularized classifier. Figure best viewed in color.

5.1 Document classification

We first test MCF predictors with blankout and Poisson corruption on document classification tasks. Specifically, we focus on three data sets: the Dmoz data set, the Reuters data set, and the Amazon review benchmark set (Blitzer et al., 2007).

Data sets. The Dmoz open directory (`http://www.dmoz.org`) contains a large collection of webpages arranged into a tree hierarchy. We use a subset consisting of $N = 8,980$ webpages from the $K = 16$ categories in the top level of the hierarchy. Each webpage is represented by a bag-of-words representation with $D = 16,498$ words. The Reuters document classification data set consists of $N = 8,293$ news articles that appeared on the Reuters newswire in 1987 belonging to $K = 65$ topics (documents corresponding to multiple topics were removed from the data). The bag-of-words representation contains $D = 18,933$ words for each document. The four Amazon data sets consist of approximately $N = 6,000$ reviews of four types of products: books, DVDs, electronics, and kitchen appliances. Each review is represented by a bag-of-words representation of the $D = 20,000$ most common words. On the Dmoz and Reuters data sets, the task is to classify the documents into one of the predefined categories. On the Amazon data set, the task is to decide whether a review is positive or negative.

Setup. On the Dmoz and Reuters data sets, we use fixed 75%/25% training/test splits. On the Amazon data set, we follow the experimental setup of Blitzer et al. (2007) by using a predefined division of the data into approximately $2,000$ training examples and about $4,000$ test examples (the exact numbers vary slightly between tasks). We perform experiments with linear classifiers that are trained using l_2-regularized quadratic, exponential, and logistic loss functions. In all experiments, the amount of l_2-regularization is determined via cross-validation. The minimization of the (expected) exponential and logistic losses is performed by running Mark Schmidt's `http://www.di.ens.fr/~mschmidt/Software/minFunc.html`minFunc-implementation of L-BFGS until convergence or until a predefined maximum number of iterations is reached. In our experiments with MCF logistic loss, we used the Jensen's upper bound of Eqn. (25).

All our predictors included a bias term that is neither regularized nor corrupted. In our experiments with MCF using blankout corruption, we use the same noise level for each feature, *i.e.* we assume that $\forall d : q_d = q$. On all data sets, we first investigate the performance of MCF as a function of the corruption level q (but we still cross-validate over the l_2-regularizer). In a second set of experiments, we cross-validate over the blankout corruption parameter q and study to what extent the performance (improvements) of MCF depend on the amount of available training data. (MCF with Poisson corruption has no additional hyperparameters, as a result of which it requires no extra cross-validations.)

Results. Figure 3 shows the test error of our MCF predictors on all data sets as a function of the blankout corruption level q. Herein, corruption level $q = 0$ corresponds to the baseline predictors, *i.e.* to predictors that do not employ MCF at all (but they do employ l_2-regularization). The results show: 1. that MCF consistently improves over standard predictors for both blankout corruption (for all corruption levels q) and Poisson corruption on five out of six tasks; 2. that MCF with Poisson corruption leads to substantial performance improvements over standard classifiers whilst introducing no additional hyperparameters; and 3. that the best performance tends to be achieved by MCF with blankout corruption

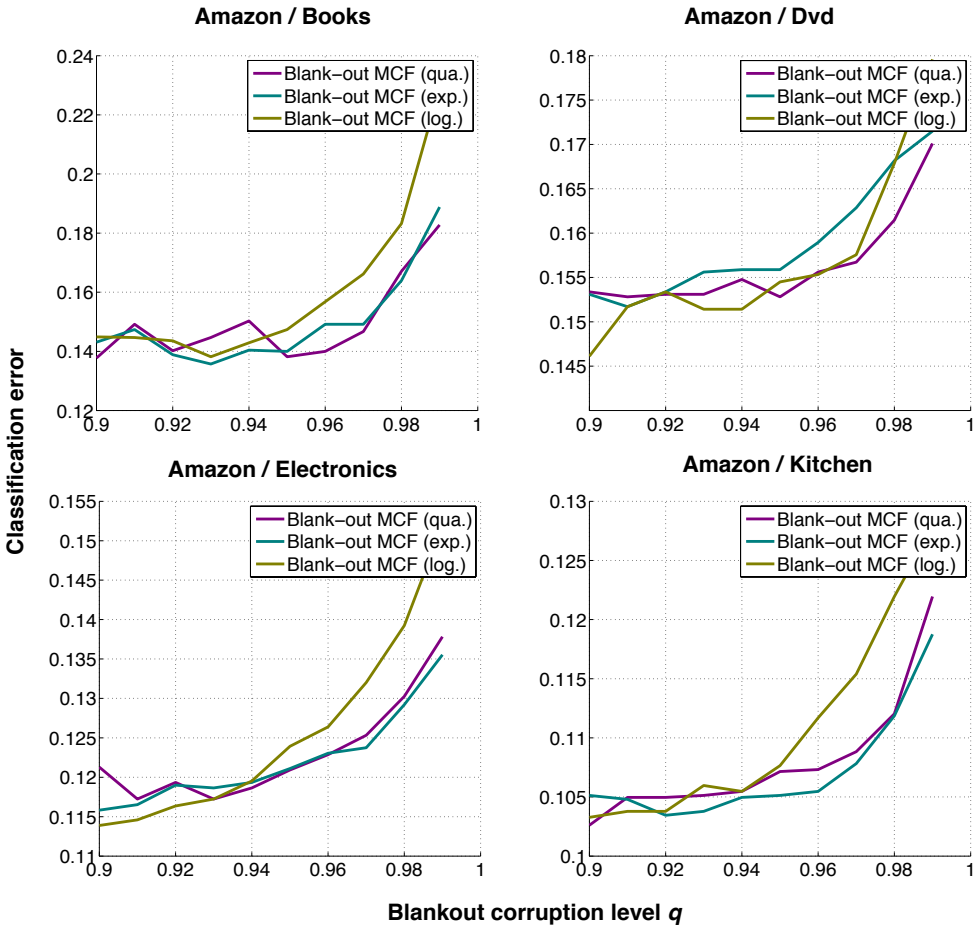

Figure 4: Classification errors of MCF predictors using blankout corruption – as a function of the blankout corruption level q in the range of $[0.9, 1]$ on the Amazon review dataset.

with relatively high corruption levels are used. In particular, the best corruption level q for the Amazon dataset is in the order of 0.8 or 0.9. Our explanation is that the ratio of relevant features (in this case, sentiment words) is fairly low in these data, and higher dropout increases the robustness of the classifiers against irrelevant features. However, further increasing the corruption level will lead to over-regularization. As shown in Figure 4, the test error of MCF predictors increases drastically as the blankout corruption level exceeds 0.9.

The best-performing MCF classifiers reduce the test errors by up to 22% on the Amazon data if q is properly set. In many of the experiments with MCF-trained losses (in particular, when blankout corruption is used), we observe that the optimal level of l_2-regularization is 0. This showcases the regularizing effect of MCF, which can render additional regularization superfluous.

Figure 5 presents the results of a second set of experiments on the Dmoz and Reuters data sets, in which we study how the performance of MCF depends on the amount of

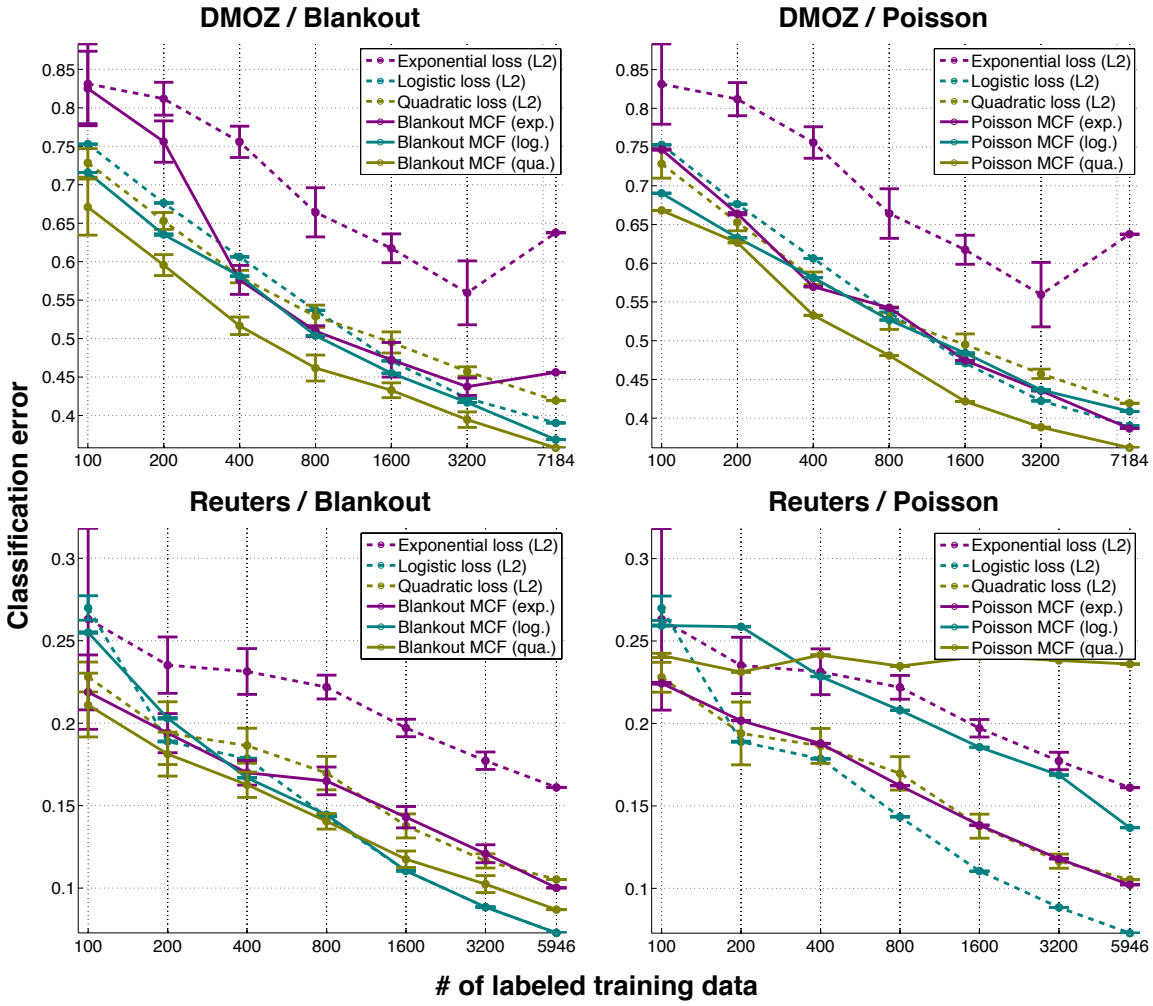

Figure 5: The performance of standard and MCF classifiers with blankout and Poisson corruption models as a function of training set size on the Dmoz and Reuters data sets. Both the standard and MCF predictors employ l_2-regularization. Figure best viewed in color.

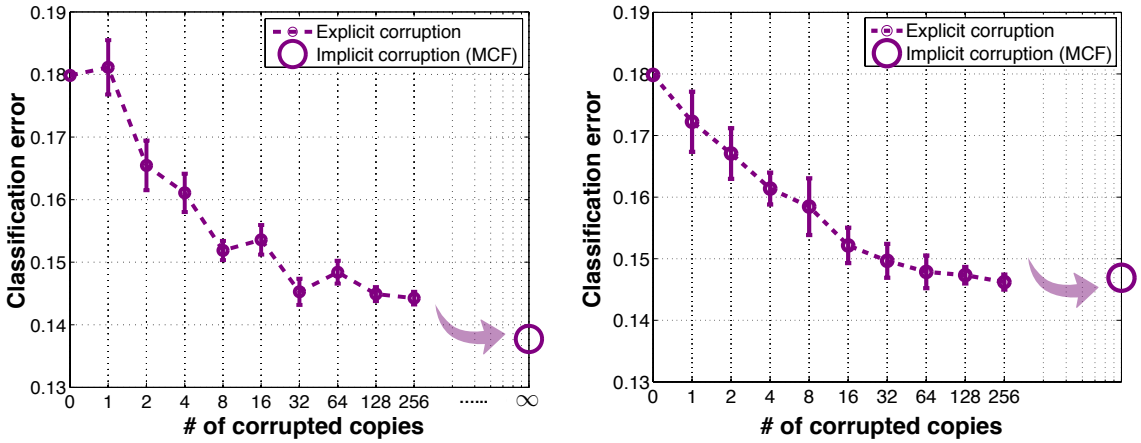

Figure 6: Comparison between MCF and explicitly adding corrupted examples to the training set (for quadratic loss) on the Amazon (books) data using blankout (*left*) and Poisson (*right* corruption. Training with MCF is equivalent to using infinitely many corrupted copies of the training data.

training data. For each training set size, we repeat the experiment five times with randomly sub-sampled training sets; the figure reports the mean test errors and the corresponding standard deviations. The results show that classifiers trained with MCF (solid curves) significantly outperform their counterparts without MCF (dashed curves). The performance improvement is consistent irrespective of the training set size, *viz.* up to 25% on the Dmoz data set.

Explicit vs. implicit feature corruption. Figure 6 shows the classification error on the Amazon Books data set when a classifier without MCF is trained on the data set with additional *explicitly* corrupted samples, as formulated in equation (3). To generate the results in the left graph of Figure 6, we use the blankout corruption model with q set by cross-validation for each training set size. To generate the right graph of Figure 6, we used the Poisson corruption model (which has no hyperparameters). In both experiments, we trained the classifiers using quadratic loss and l_2-regularization (with cross-validation over the regularization parameter for each data size).

The graphs in the figure show a clear trend that the error *decreases* when the training set contains more corrupted versions of the original training data, *i.e.* with higher M in equation (3). The graphs also illustrate that the best performance is obtained as M approaches infinity, which is equivalent to MCF with blankout or Poisson corruption (big markers in the bottom right). This nicely illustrates the potential of MCF to improve generalization by training on corrupted examples without increasing the computational complexity of learning.

Comparing approximations and bounds for MCF logistic loss. We also performed experiments to evaluate the effectiveness of the three approaches to deal with MCF logistic loss, *viz.* the quadratic approximation of Eqn. (20), the Jensen's upper bound of Eqn. (25), and the variational upper bound of Eqn. (28). Specifically, we compared the performance of the three approaches when using blankout corruption with different values

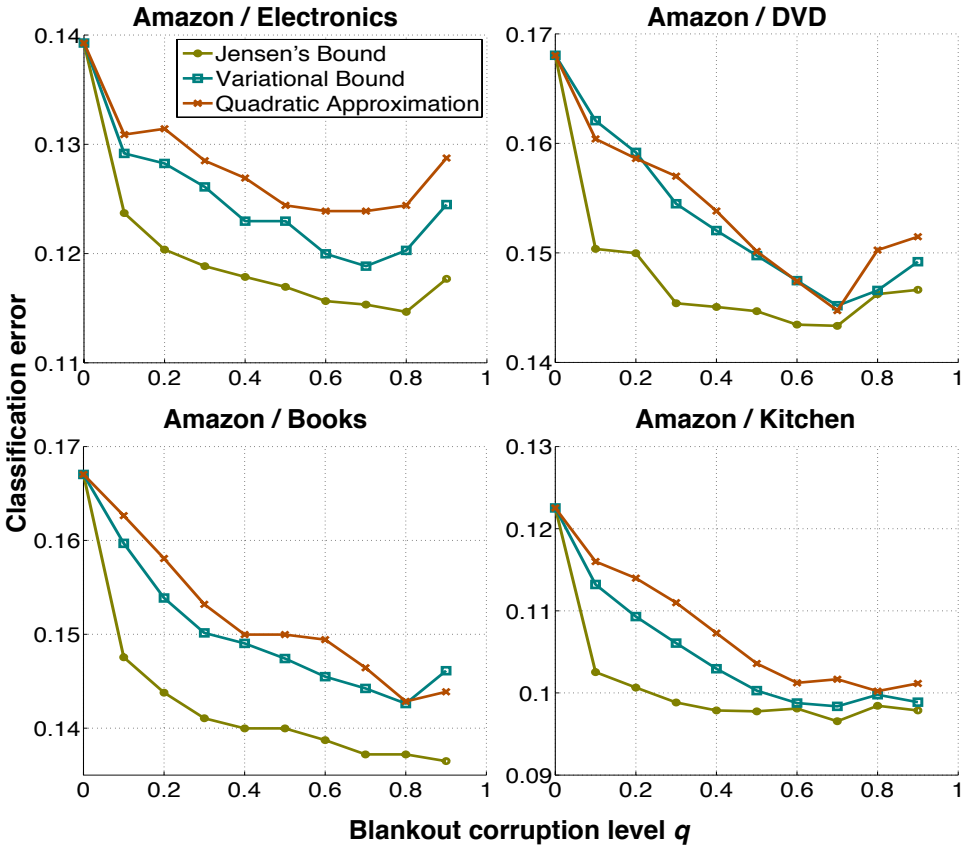

Figure 7: Comparison between the three approaches (quadratic approximation, Jensen's upper bound, and variational upper bound) to MCF logistic loss with blankout as a function of the blankout rate q. The value $q = 0$ corresponds to classifiers trained using standard logistic loss. Figure best viewed in color.

of q on the four Amazon data sets. As before, all classifiers used additional l_2-regularization; the regularization parameters were set by cross-validating on a held-out validation set.

The results of these experiments are shown in Figure 7 (please note that, again, the setting $q = 0$ corresponds to training classifiers using standard logistic loss). The results presented in the figure show that classifiers trained using logistic loss with MCF-blankout outperform classifiers that are trained using standard logistic loss, irrespective of whether the quadratic approximation or one of the two upper bounds is used. Having said that, the results also show that the two upper bounds on MCF logistic loss outperform the quadratic approximation in most experiments. In addition, the results show that the use of Jensen's upper bound tends to lead to better results than the use of the variational upper bound to the logistic loss, despite the fact that the variational bound is always at least as tight as Jensen's bound. We surmise this results is due to the fact that the variational bound of Eqn. (28) is not jointly convex in $\mathbf{w}$ and λ, whereas the Jensen's bound of Eqn. (25) is

	Quadratic	Exponential	Logistic
No MCF	32.6%	39.7%	32.5%
Poisson MCF	29.1%	39.5%	30.0%
Blankout MCF	32.3%	37.9%	29.4%

Table 3: Classification errors obtained on the CIFAR-10 data set with MCF classifiers trained on simple spatial-pyramid bag-of-visual-words features.

convex in $\mathbf{w}$. As a result, we may get stuck in poorer local minima of the loss when using the variational upper bound.

5.2 Image classification

We perform image-classification experiments with MCF on the CIFAR-10 data set (Krizhevsky, 2009), which is a subset of the 80 million tiny images (Torralba et al., 2008). The data set contains RGB images with 10 classes of size 32×32, and consists of a fixed training set of $50,000$ images and a fixed test set of $10,000$ images.

Setup. We follow the experimental setup of Coates et al. (2011): we whiten the training images and extract a set of 7×7 image patches on which we apply k-means clustering (with $k = 2048$) to construct a codebook of prototypical image patches. Next, we slide a 7×7 pixel window over each image and identify the nearest prototype in the codebook for each window location. We construct an image descriptor by subdividing the image into four equally sized quadrants and counting the number of times each prototype occurs in each quadrant, which leads to a descriptor of dimensionality $D = 4 \times 2048$. This way of extracting the image features is referred to by Coates et al. (2011) as k-means with hard assignment, average pooling, patch size 7×7, and stride 1. Because all images have the same size, we do not need to normalize the descriptors. We train MCF predictors with blankout and Poisson corruption on the full set of training images, cross-validating over a range of l_2-regularization parameters. The generalization error is evaluated on the test set. For our experiments with logistic loss, we used the Jensen's upper bound of Eqn. (25).

Results. The results are reported in Table 3. The baseline classifiers (without MCF) are comparable to the 68.8% accuracy reported by Coates et al. (2011) with exactly the same experimental setup (except for exponential loss). The results illustrate the potential of MCF classifiers to improve the prediction performance on bag-of-visual-words features, in particular, when using quadratic or logistic loss in combination with Poisson corruption.

Although our focus in this section is to merely illustrate the potential of MCF on image classification tasks, it is worth noting that the best results in Table 3 match those of a highly non-linear mean-covariance RBMs trained on the same data (Ranzato and Hinton, 2010), despite our use of very simple visual features and of linear classifiers.

5.3 Analysis of DNA

We also performed experiments with MCF in which the aim is to automatically recognize splice-junctions in DNA: the locations that indicate the boundaries between regions of

DNA that are used in the production of mRNA (so-called introns) and regions of DNA that are removed in the production of mRNA (exons). The two types of boundaries are called *acceptors* (exon-to-intron boundaries) and *donors* (intron-to-exon boundaries). The mRNA that is created using the information in the DNA is subsequently translated into proteins, which form the basis of all chemical processes inside the cell. Therefore, automatic recognition of splice-junctions is essential in the analysis of cell processes.

Setup. We perform experiments on the data set constructed by Noordewier et al. (1991). The data set comprises 3190 examples of three classes: it contains 767 examples of acceptors, 768 examples of donors, and 1655 negative examples. Each example is represented by a DNA subsequence of 120 nucleotides: 60 on each side of the boundary for positive examples (acceptors and donors). We represent the nucleotides using a 1-of-K representation, where $K = 4$ is the number of distinct nucleotides (A, G, T, and C). This leads to binary feature vectors of dimensionality $D = 480$.

We measure the performance of our classifiers using 10-fold cross-validation. We repeat each experiment fifty times and average the performances across the fifty runs. In this experiment, we focus on MCF with logistic loss because in preliminary experiments, models trained with logistic loss produced substantially better results than models that used other loss functions. As in the previous experiments, we use the Jensen's upper bound on MCF-logistic loss of Eqn. (25) as surrogate loss function. Because of the binary nature of our data, we focus on using blankout corruption in MCF. In addition, we explore a simple *multinomial* corruption model that aims to model the sequencing errors that are made by DNA sequencing machines (*e.g.*, Le et al. (2013); Zagordi et al. (2010)). The multinomial corruption model assumes that a base pair $x \in \{1, \ldots, K\}$ is an erroneous read with probability q, and assumes a uniform distribution over all alternative base pairs in case of an erroneous read:

$$p(\tilde{x} = x) = 1 - q \text{ and } \forall k \neq x : p(\tilde{x} = k) = \frac{q}{K - 1}. \tag{31}$$

We note here that this corrupting distribution violates the assumption that $\mathbb{E}[\tilde{\mathbf{x}}] = \mathbf{x}$. However, this is not a problem since we are focusing on a classification task (the unbiasedness assumption is only relevant to regression tasks).

Results. Figure 8 presents the performance of our MCF classifiers with blankout corruption (left) and multinomial corruption (right) as a function of the MCF corruption rate q. The results presented in the figure show that MCF with blankout may improve the performance of linear models for the recognition of splice-junctions by more than 0.6%, which corresponds to a relative improvement of approximately 15%. Improvements in performance are obtained for a wide range of blankout rates, with optimal blankout rates ranging between 0.15 and 0.4. The results also show that the multinomial corruption model is too simple to lead to performance improvements. However, we expect that these results can be substantially improved by incorporating in the corruption process knowledge about: 1. the *coverage* obtained on a particular base pair (base-pair reads with higher coverage are read in more DNA snippets and are, therefore, less likely to be erroneous) and 2. knowledge on the error distribution of a DNA sequencing machine (for instance, sequencing machines are more likely to confuse A and T than A and G). Both types of knowledge are available for modern DNA sequencing processes (Le et al., 2013; Zagordi et al., 2010). The MCF frame-

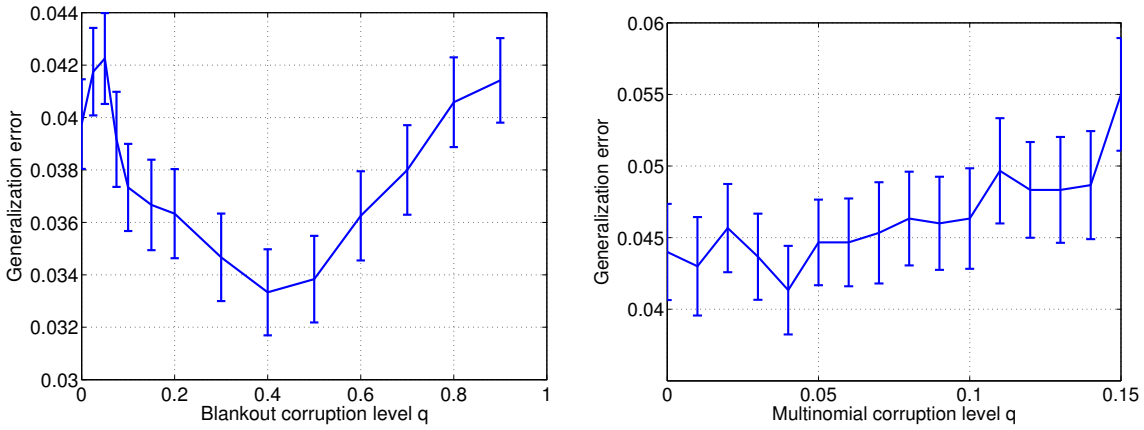

Figure 8: Generalization error of splice-junction classifiers trained using MCF-logistic loss with blankout corruption (left) and multinomial corruption (right) as a function of the corruption rate q (lower is better). All results are obtained by averaging over 50 runs of 10-fold cross-validation. The error bars indicate the standard error of the generalization error estimate.

work allows us to model such domain knowledge into the the corruption model. However, since no information on the coverage and sequencing error distribution is available for our data set (Noordewier et al., 1991), we leave such extensions of the DNA experiments to future work.

5.4 Nightmare at test time

To test the performance of our MCF predictors with blankout corruption under the "nightmare at test time" scenario, we perform experiments on the MNIST data set of handwritten digit images. The MNIST data set contains $N = 60,000$ training and $10,000$ test images of size $D = 28 \times 28 = 784$ pixels, and comprises $K = 10$ classes.

Setup. We train our predictors on the full training set, and evaluate their performance on versions of the test set in which a certain percentage of the pixels are randomly blanked out, *i.e.* set to zero. We compare the performance of our MCF-predictors (using blankout corruption) with that of standard predictors that use l_1 or l_2-regularized quadratic, exponential, logistic, and hinge loss. As before, we used the Jensen's upper bound for predictors with MCF logistic loss, and we use cross-validation to determine the optimal value of the regularization parameter. For MCF predictors, we also cross-validate over the blankout corruption level q (again, we use the same noise level for each feature, *i.e.* $\forall d : q_d = q$). In addition to the comparisons with standard predictors, we also compare the performance of MCF with that of FDROP (Globerson and Roweis, 2006), which is a state-of-the-art algorithm for the "nightmare at test time" setting that minimizes the hinge loss under an adversarial worst-case scenario (see section 2).

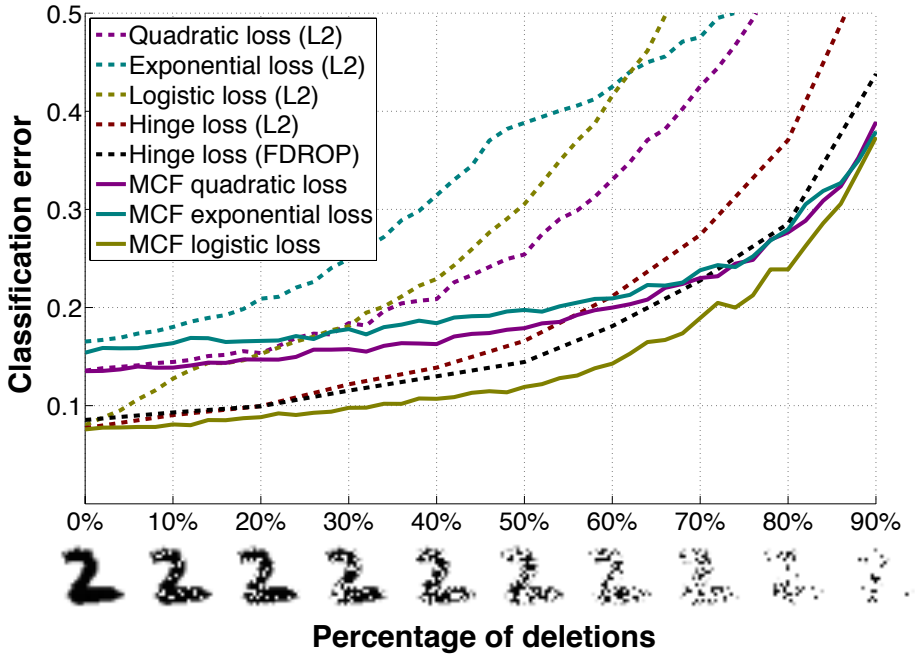

Figure 9: Evaluation on the "nightmare at test-time" scenario. Classification errors of standard and MCF predictors with a blankout corruption model – trained using three different losses – and of FDROP (Globerson and Roweis, 2006) on the MNIST data set using the "nightmare at test time" scenario. Classification errors are represented on the y-axis, whereas the amount of features that are deleted out at test time is represented on the x-axis. The images of the digit illustrate the amount of feature deletions applied on the digit images that are used as test data. Figure best viewed in color.

The performances are reported as a function of the feature-deletion percentage in the test set, *i.e.* as a function of the probability with which a pixel in the test set is switched off. Following the experimental setting of Globerson and Roweis (2006), we perform the cross-validation for each deletion percentage independently, *i.e.* we create a small validation set with the same feature-deletion level and use it to determine the best regularization parameters and blankout corruption level q for that percentage of feature deletions. We are thus assuming a transductive learning setting, which is a common assumption in many learning scenarios with domain shift (Blitzer et al., 2007; Daumé III, 2007).

Results. Figure 9 shows the performance of our predictors as a function of the percentage of deletions in the test images. The figure shows the performance for all three loss functions with MCF (solid lines) and without MCF (dashed lines). The performance of a standard predictor using hinge loss is shown as a red dashed line; the performance of FDROP (a state-of-the-art method for this learning setting) is shown as a black dashed line. The results presented in Figure 9 clearly illustrate the ability of MCF with blankout corruption to produce predictors that are more robust to the "nightmare at test time" scenario: MCF

improves the performance substantially for all three loss functions considered. For instance, in the case in which 50% of the pixels in the test images is deleted, the performance improvements obtained using MCF for quadratic, exponential, and logistic loss are 40%, 47%, and 60%, respectively. Further, the results also indicate that MCF-losses may outperform FDROP; in particular, our average-case logistic loss outperforms FDROP's worst-case hinge loss across the board[3]. This is particularly impressive as FDROP uses the hinge loss, which performs surprisingly better than the alternative losses on this data set (the improvement of FDROP over the generic hinge-loss is in fact relatively modest). This result suggests that it is better to consider an average-case scenario than a worst-case scenario in the "nightmare at test time" setting.

6. Discussion and Conclusion

We presented an approach to learn classifiers by marginalizing corrupted features (MCF). Specifically, MCF trains predictors by introducing corruption on the training examples, which is marginalized out in the expectation of the loss function. We minimize the expected loss with respect to the model parameters. Our experimental results show that MCF predictors with blankout and Poisson corruption perform very well in the context of bag-of-words features, whilst MCF with blankout corruption was also very effective on DNA data. MCF with blankout corruption appears to be most effective in high-dimensional learning problems in which the data is sparse. MCF with Poisson corruption is particularly interesting for count features, as it improves classification performances without introducing any additional hyper-parameters. As a disclaimer, care must be taken when applying MCF with Poisson corruption on data sets with outliers. The Poisson corruption may strongly emphasize outliers in the expected loss because the variance of a Poisson distribution is identical to its mean, and because we use loss functions that are not robust to outliers. A simple solution to this problem may be to redefine the corruption distribution to $p(\tilde{x}_d|x_d) = Pois(\tilde{x}_{nd}|\min\{x_{nd}, u\})$ for some cutoff parameter $u \geq 0$.

Throughout our experiments with MCF on exponential and logistic losses, the weight of the l_2-regularization term, which was set by cross-validation, typically ended up very close to zero. This implies that the regularizing effect of MCF corruption is sufficient even for high dimensional data. In comparison with traditional regularization, MCF corruption shines in two ways: 1. it often yields superior classification performance and 2. it can be much more intuitive to set parameters about data corruption than about model hyper parameters. Further, MCF with blankout corruption also appears to prevent *weight undertraining* (Sutton et al., 2005): it encourages the weight on each feature to be non-zero, in case this particular feature survives the corruption.

The resulting redundancy is one of the key factors why MCF corruption makes classifiers so effective against the "nightmare at test-time" scenario. Consequently, MCF classifiers are particularly useful in learning settings in which features in the test data may be missing (*e.g.*, because sensors that were measuring these features temporarily broke down). Learning with MCF is quite different from previous approaches for this setting (Dekel and Shamir, 2008; Globerson and Roweis, 2006). In particular, prior work focuses on the worst-case sce-

3. Quadratic and exponential losses perform somewhat worse because they are less appropriate for linear classifiers, but even they outperform FDROP for large numbers of feature deletions in the test data.

nario whereas we explicitly consider the (arguably) more common average-case scenario by considering all possible corrupted observations. This has the advantage that it is computationally much cheaper and that it allows for incorporating prior knowledge in the learning. For instance, if the data is generated by a collection of unreliable sensors, knowledge on the reliability of a sensor may be used to set the corresponding q_d-parameter. The strong performance of MCF in the "nightmare at test-time" scenario suggest it be well suited for learning under domain shift (Blitzer et al., 2007; Daumé III, 2007). In particular, as the corrupting distribution may be used to shift the data distribution in the source domain towards the data distribution in the target domain — potentially, by learning the parameters of the corrupting distribution using maximum likelihood. We leave such investigations for future work.

Another interesting direction for future work is to investigate extensions of MCF to structured prediction (Wang et al., 2013), as well as to investigate if MCF can be employed for kernel machines. We also plan to explore in more detail what corruption models $p(\tilde{\mathbf{x}}|\mathbf{x})$ are useful for what types of data. Further, MCF could be used in the training of neural networks with a single layer of hidden units: blankout noise *on the hidden nodes* can improve the performance of the networks (LeCun et al., 1990; Hinton et al., 2012) and can be marginalized out analytically. Moreover, classifiers trained using MCF-blankout may be very well suited for application in *anytime classification* scenarios (Grubb and Bagnell, 2012), as they are optimized to work well on all possible subsets of features. Another interesting direction for future work is to investigate the theoretical properties of learning with MCF. Xu et al. (2009) have derived generalization bounds for learning with corruption under the worst-case scenario (*i.e.* for the related models discussed in Section 2), and McAllester (2013) recently derived a PAC-Bayesian bound for learning linear models with blankout in the average-case scenario; however, generic bounds for learning with MCF do not yet exist. A final interesting direction would be to investigate the effect of marginalizing *corrupted labels* or target values (Lawrence and Schölkopf, 2001; Chen et al., 2013).

Acknowledgements

The authors thank Fei Sha, Wouter Kouw, Max Welling, and Lawrence Saul for helpful discussions. In particular, the variational upper bound for the logistic loss was suggested by Lawrence Saul. Laurens van der Maaten acknowledges support from the Netherlands Organisation for Scientific Research (NWO; grant 612.001.301), from the EU-FP7 SSPNet and INSIDDE projects, and from the AAL project SALIG++. Kilian Weinberger, Minmin Chen, and Stephen Tyree are supported by NSF grants 1149882 and 1137211.

Appendix A. Overview of loss functions

Table 4 gives an overview of all MCF loss functions discussed in the paper. Table 5 lists the associated gradients required for learning. For an overview of moment-generating functions and their gradients, we refer to Table 2.

4. For MCF quadratic loss, the optimal parameters can be identified via the closed-form solution in (6).

MCF Loss	y	Loss function
Quadratic	$\{-1,+1\}$	$\mathbf{w}^\top \left(\mathbb{E}[\tilde{\mathbf{x}}_n]\mathbb{E}[\tilde{\mathbf{x}}_n]^\top + \mathbb{V}[\tilde{\mathbf{x}}_n] \right) \mathbf{w} - 2 \left(y_n \mathbb{E}\left[\tilde{\mathbf{x}}_n\right] \right)^\top \mathbf{w} + 1$
Exponential	$\{0,1\}$	$\prod_{d=1}^{D} \mathbb{E}\left[\exp\left(-y_n w_d \tilde{x}_{nd}\right)\right]$
Log.–Quadr.	$\{-1,+1\}$	$-y\mathbf{w}^\top\mathbf{x} + A(\mathbf{w}^\top\mathbf{x}) + \frac{1}{2}\frac{\partial^2 A(\mathbf{w}^\top\mathbf{x})}{\partial(\mathbf{w}^\top\mathbf{x})^2}\mathbb{V}[\mathbf{w}^\top\tilde{\mathbf{x}}_n]$
Log.–Jensen	$\{0,1\}$	$\log\left(1 + \prod_{d=1}^{D} \mathbb{E}\left[\exp\left(-y_n w_d \tilde{x}_{nd}\right)\right]\right)$
Log.–Variat.	$\{0,1\}$	$\min_{\lambda\in[0,1]} \left[\log\left(\prod_{d=1}^{D} \mathbb{E}[(\exp(\lambda y_n w_d \tilde{x}_{nd})] + \prod_{d=1}^{D} \mathbb{E}[\exp((\lambda-1) y_n w_d \tilde{x}_{nd})] \right) - \lambda y_n \mathbf{w}^\top \mathbf{x}_n \right]$

Table 4: MCF losses on a labeled instance $(\mathbf{x}_n, y_n)$. For an overview of moment-generating functions, please refer to Table 2. All the expectations and variances are under the corrupting distribution $p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)$.

MCF Loss	Loss gradient
Quadratic	Not applicable[4]
Exponential	$\frac{\partial \mathbb{E}[\exp(-y_n w_d \tilde{x}_{nd})]}{\partial w_d} \prod_{d' \neq d} \mathbb{E}\left[\exp\left(-y_n w_{d'} \tilde{x}_{nd'}\right)\right]$
Log.–Quadr.	$-y_n\mathbf{x}_n + 2\sigma(\mathbf{w}^\top\mathbf{x}_n) - 1 + \frac{\partial^2 A(\mathbf{w}^\top\mathbf{x}_n)}{\partial(\mathbf{w}^\top\mathbf{x}_n)^2}(\mathbb{E}[\tilde{\mathbf{x}}_n\tilde{\mathbf{x}}_n^\top] - \mathbf{x}_n\mathbf{x}_n^\top)\mathbf{w} + \frac{1}{2}\frac{\partial^3 A(\mathbf{w}^\top\mathbf{x}_n)}{\partial(\mathbf{w}^\top\mathbf{x}_n)^3}\mathbb{V}[\mathbf{w}^\top\tilde{\mathbf{x}}_n]$
Log.–Jensen	$\frac{\prod_{d'\neq d}\mathbb{E}[\exp(-y_n w_{d'}\tilde{x}_{nd'})]}{1+\prod_{d'=1}^{D}\mathbb{E}[\exp(-y_n w_{d'}\tilde{x}_{nd'})]} \frac{\partial \mathbb{E}[\exp(-y_n w_d \tilde{x}_{nd})]}{\partial w_d}$
Log.–Variat.	$-\lambda^* y_n x_{nd} + \frac{\exp(\lambda^*)\prod_{d'\neq d}\mathbb{E}[\exp(-y_n w_{d'}\tilde{x}_{nd'})]}{\prod_{d'=1}^{D}\mathbb{E}[\exp(\lambda^* y_n w_{d'}\tilde{x}_{nd'})] + \prod_{d'=1}^{D}\mathbb{E}[\exp((\lambda^*-1) y_n w_{d'}\tilde{x}_{nd'})]} \frac{\partial \mathbb{E}[\exp(-y_n w_d \tilde{x}_{nd})]}{\partial w_d}$ $\left[\prod_{d'=1}^{D}\mathbb{E}\left[\exp\left((\lambda^*-1) y_n w_{d'}\tilde{x}_{nd'}\right)\right] + \exp(-1)\prod_{d'=1}^{D}\mathbb{E}\left[\exp\left((\lambda^*-2) y_n w_{d'}\tilde{x}_{nd'}\right)\right]\right]$

Table 5: Derivatives required for learning from a labeled instance $(\mathbf{x}_n, y_n)$ using various MCF losses. For an overview of the derivatives of moment-generating functions, please refer to Table 2. All the expectations are under the corrupting distribution $p(\tilde{\mathbf{x}}_n|\mathbf{x}_n)$. In the variational upper bound of the logistic loss, λ^* is obtained by performing the minimization described in Table 4.

Appendix B. Numerical stability

When using dropout, the Jensen's upper bound on the expected logistic loss is not numerically stable. In particular, the upper bound $\hat{\mathcal{L}}$ takes the following form for multi-class classification:

$$\hat{\mathcal{L}}(\mathcal{D}; \mathbf{w}) = -\sum_{n=1}^{N} \log\left[\frac{\prod_{d=1}^{D}(q_d + (1-q_d)\exp(\mathbf{y}_n^\top \mathbf{w}_d x_{nd}))}{\sum_{\mathbf{y'}_n}\prod_{d=1}^{D}(q_d + (1-q_d)\exp(\mathbf{y'}_n^\top \mathbf{w}_d x_{nd}))} \right]. \tag{32}$$

As is common in logistic regression, we can evaluate this expression by expressing the nominator and denominator in the log-domain, performing a log-shift, and exponentiating both sides:

$$\hat{\mathcal{L}}(\mathcal{D};\mathbf{w}) = -\sum_{n=1}^{N} \log \left[\frac{\exp\left(s + \sum_{d=1}^{D} \log\left[q_d + (1-q_d)\exp(\mathbf{y}_n^\top \mathbf{w}_d x_{nd})\right]\right)}{\sum_{\mathbf{y'}_n} \exp\left(s + \sum_{d=1}^{D} \log\left[q_d + (1-q_d)\exp(\mathbf{y'}_n^\top \mathbf{w}_d x_{nd})\right]\right)} \right], \quad (33)$$

for some appropriately chosen log-shift s. However, the computation of this quantity may be problematic when $\mathbf{y}_n^\top \mathbf{w}_d x_{nd} > 0$. We resolve this by noting the following equality:

$$\log\left[q_d + (1-q_d)\exp(\mathbf{y}_n^\top \mathbf{w}_d x_{nd})\right] = \mathbf{y}_n^\top \mathbf{w}_d x_{nd} + \log\left[q_d \exp(-\mathbf{y}_n^\top \mathbf{w}_d x_{nd}) + 1 - q_d\right]. \quad (34)$$

Depending on the sign of $\mathbf{y}_n^\top \mathbf{w}_d x_{nd}$, we compute the left-hand side or the right-hand side of this equation.

References

J. Abernethy, E. Hazan, and A. Rakhlin. Competing in the dark: An efficient algorithm for bandit linear optimization. In *Proceedings of the Conference on Learning Theory*, pages 263–274, 2008.

D.M. Allen. The relationship between variable selection and data augmentation and a method for prediction. *Technometrics*, 16:125–127, 1974.

C. Bhattacharyya, K.S. Pannagadatta, and A.J. Smola. A second order cone programming formulation for classifying missing data. In *Advances in Neural Information Processing Systems*, pages 153–160, 2004.

J. Blitzer, M. Dredze, and F. Pereira. Biographies, Bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Association for Computational Linguistics*, volume 45, page 440, 2007.

S.P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, New York, NY, 2004.

C.J.C. Burges and B. Schölkopf. Improving the accuracy and speed of support vector machines. *Advances in Neural Information Processing Systems*, 9:375–381, 1997.

N. Cesa-Bianchi, S. Shalev-Shwartz, and O. Shamir. Online learning of noisy data with kernels. In *Proceedings of the Conference on Learning Theory*, pages 218–230, 2010.

N. Cesa-Bianchi, S. Shalev-Shwartz, and O. Shamir. Online learning of noisy data. *IEEE Transactions on Information Theory*, 57(12):7907–7931, 2011.

O. Chapelle, J. Weston, L. Bottou, and V. Vapnik. Vicinal risk minimization. In *Advances in Neural Information Processing Systems*, pages 416–422, 2000.

G. Chechik, G. Heitz, G. Elidan, P. Abbeel, and D. Koller. Max-margin classification of data with absent features. *Journal of Machine Learning Research*, 9(Jan):1–21, 2008.

M. Chen, Z. Xu, K.Q. Weinberger, and F. Sha. Marginalized denoising autoencoders for domain adaptation. In *Proceedings of the International Conference on Machine Learning*, pages 767–774, 2012.

M. Chen, A. Zheng, and K.Q. Weinberger. Fast image tagging. In *Proceedings of* 30^{th} *International Conference on Machine Learning*, pages 1274–1282, 2013.

A. Coates, H. Lee, and A.Y. Ng. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the International Conference on Artificial Intelligence & Statistics, JMLR W&CP 15*, pages 215–223, 2011.

T. Cover and P. Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13(1):21–27, 1967.

H. Daumé III. Frustratingly easy domain adaptation. In *Proceedings of the* 45^{th} *Annual Meeting of the Association of Computational Linguistics*, pages 256–263, 2007.

O. Dekel and O. Shamir. Learning to classify with missing and corrupted features. In *Proceedings of the International Conference on Machine Learning*, pages 216–223, 2008.

R.O. Duda, P.E. Hart, and D.G. Stork. *Pattern Classification.* Wiley Interscience Inc., 2001.

A. Flaxman, A. Kalai, and H. McMahan. Online convex optimization in the bandit setting: Gradient descent without a gradient. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 385–394, 2005.

Y. Freund and R.E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In *Proceedings of the Second European Conference on Computational Learning Theory*, pages 23–37, 1995.

A. Globerson and S. Roweis. Nightmare at test time: Robust learning by feature deletion. In *Proceedings of the International Conference on Machine Learning*, pages 353–360, 2006.

X. Glorot, A. Bordes, and Y. Bengio. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *Proceedings of the International Conference on Machine Learning*, pages 513–520, 2011.

E. Grave, G. Obozinski, and F. Bach. Trace lasso: A trace norm regularization for correlated designs. In *Advances in Neural Information Processing Systems 24*, pages 2187–2195, 2011.

A. Grubb and J.A. Bagnell. Speedboost: Anytime prediction with uniform near-optimality. In *Proceedings of the International Conference on Artificial Intelligence & Statistics (AI-STATS)*, pages 458–466, 2012.

R. Herbrich and T. Graepel. Invariant pattern recognition by semidefinite programming machines. In *Advances in Neural Information Processing Systems*, volume 16, page 33, 2004.

G.E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R.R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors, 2012.

A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

N.D. Lawrence and B. Schölkopf. Estimating a kernel Fisher discriminant in the presence of label noise. In *Proceedings of the International Conference in Machine Learning*, pages 306–313, 2001.

H.S. Le, M.H. Schulz, B.M. McCauley, V.F. Hinman, and Z. Bar-Joseph. Probabilistic error correction for RNA sequencing. *Nucleic Acids Research*, 41(10):e109, 2013.

Y. LeCun, J.S. Denker, and S.A. Solla. Optimal brain damage. In *Advances in Neural Information Processing Systems*, pages 598–605, 1990.

R. Livni, K. Crammer, and A. Globerson. A simple geometric interpretation of svm using stochastic adversaries. In *Proceedings of the* 15^{th} *International Conference on Artificial Intelligence and Statistics (AI-STATS)*, pages 722–730, 2012.

F. Maillet, D. Eck, G. Desjardins, and P. Lamere. Steerable playlist generation by learning song similarity from radio station playlists. In *Proceedings of the International Society for Music Information Retrieval Conference*, pages 345–350, 2009.

D. McAllester. A PAC-Bayesian tutorial with a dropout bound, 2013.

G. Mesnil, Y. Dauphin, X. Glorot, S. Rifai, Y. Bengio, I. Goodfellow, E. Lavoie, X. Muller, G. Desjardins, D. Warde-Farley, P. Vincent, A. Courville, and J. Bergstra. Unsupervised and transfer learning challenge: a deep learning approach. *JMLR: Workshop and Conference Proceedings*, 7:1–15, 2011.

M.O. Noordewier, G.G. Towell, and J.W. Shavlik. Training knowledge-based neural networks to recognize genes in DNA sequences. In *Advances in Neural Information Processing Systems*, volume 3, pages 530–536, 1991.

E. Parzen. On estimation of a probability density function and mode. *Annals of Mathematical Statistics*, 33:1065–1076, 1962.

M. Ranzato and G.E. Hinton. Modeling pixel means and covariances using factorized third-order Boltzmann machines. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2551–2558, 2010.

A. Rostamizadeh, A. Agarwal, and P. Bartlett. Learning with missing features. In *Proceedings of Uncertainty in Artificial Intelligence*, pages 635–642, 2011.

F. Sha and F. Pereira. Shallow parsing with Conditional Random Fields. In *Proceedings of Human Language Technology – NAACL 2003*, pages 213–220, 2003.

G. Shafer and V. Vovk. A tutorial on conformal prediction. *The Journal of Machine Learning Research*, 9:371–421, 2008.

P.K. Shivaswamy, C. Bhattacharyya, and A.J. Smola. Second order cone programming approaches for handling missing and uncertain data. *Journal of Machine Learning Research*, 7:1283–1314, 2006.

C. Sutton, M. Sindelar, and A. McCallum. Feature bagging: Preventing weight undertraining in structured discriminative learning. Technical Report IR-402, University of Massachusetts, 2005.

C.H. Teo, A. Globerson, S. Roweis, and A. Smola. Convex learning with invariances. *Advances in Neural Information Processing Systems*, 20:1489–1496, 2008.

A. Torralba, R. Fergus, and W.T. Freeman. 80 million tiny images: A large dataset for non-parametric object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(11):1958–1970, 2008.

T. Trafalis and R. Gilbert. Robust support vector machines for classification and computational issues. *Optimization Methods and Software*, 22(1):187–198, 2007.

L.J.P. van der Maaten, M. Chen, S. Tyree, and K.Q. Weinberger. Learning by marginalizing corrupted features. In *Proceedings of the International Conference on Machine Learning*, pages 410–418, 2013.

P. Vincent, H. Larochelle, Y. Bengio, and P.A. Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the International Conference on Machine Learning*, pages 1096–1103, 2008.

P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11(Dec):3371–3408, 2010.

S. Wager, S. Wang, and P. Liang. Dropout training as adaptive regularization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 351–359, 2013.

S. Wang, M. Wang, S. Wager, P. Liang, and C.D. Manning. Feature noising for log-linear structured prediction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1170–1179, 2013.

H. Xu, C. Caramanis, and S. Mannor. Robustness and regularization of support vector machines. *Journal of Machine Learning Research*, 10:1485–1510, 2009.

O. Zagordi, R. Klein, M. Däumer, and N. Beerenwinkel. Error correction of next-generation sequencing data and reliable estimation of HIV quasispecies. *Nucleic Acids Research*, 38 (21):7400–7409, 2010.

J. Zhu, S. Rosset, H. Zou, and T. Hastie. Multi-class AdaBoost. Technical Report 430, Department of Statistics, University of Michigan, 2006.